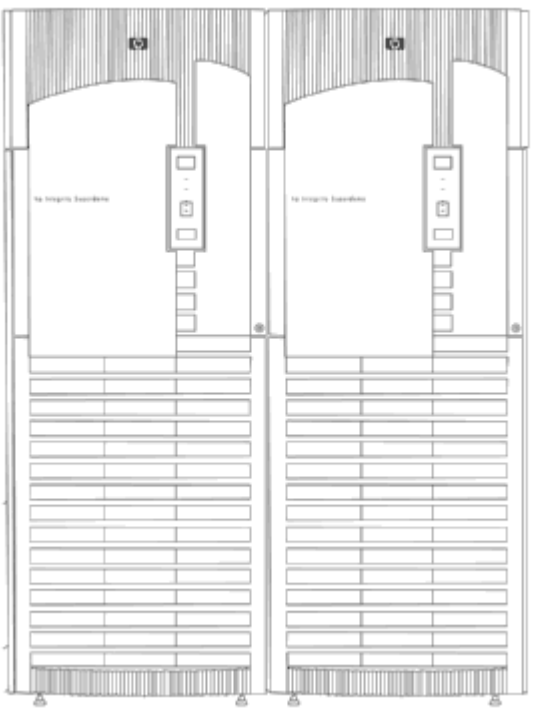

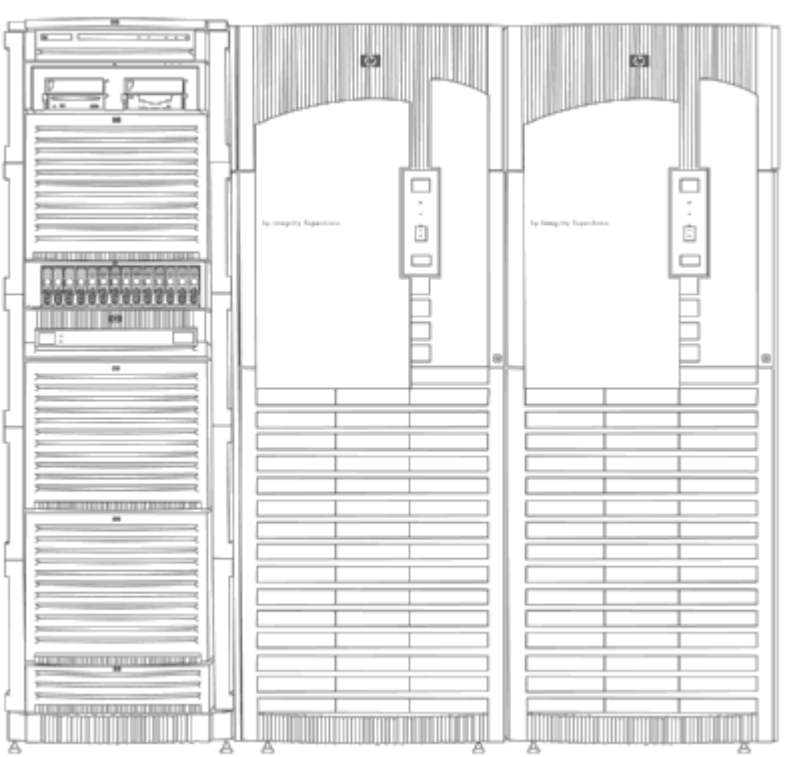

Overview

At A Glance

The latest release of Superdome, HP Integrity Superdome supports the new and improved sx2000 chip set. The Integrity Superdome with the sx2000 chipset supports the Mad9M Itanium 2 processor at initial release. HP Integrity Superdome supports the following processors:

sx2000

- Dual core Itanium 2 1.6-GHz processor
- Itanium 2 1.6-GHz processor

sx1000

- Itanium 2 1.5 GHz and 1.6 GHz processors
- HP mx2 processor module based on two Itanium 2 processors

HP sx1000 Integrity Superdome supports mixing the Itanium 2 1.5 GHz processor, the Itanium 2 1.6 GHz processor and the HP mx2 processor module in the same system, but on different partitions, as long as they have the same chipset. All cell boards to be mixed in a system must contain the same chipset, sx1000 or sx2000, but not both. HP Integrity sx1000 Superdome also supports mixing the Itanium 2 1.6 GHz processor, PA 8800 and PA 8900 processors in the same system, but on different partitions, again only with the same chipset, sx1000 or sx2000.

Throughout the rest of this document, the term HP sx1000 Integrity Superdome with Itanium 2 1.5 GHz processors, Itanium 2 1.6 GHz processors or mx2 processor modules will be referred to as simply "Superdome sx1000". The HP sx2000 Integrity Superdome with dual core Itanium processors or single core Itanium 2 processors will be referred to as "Superdome sx2000".

Superdome with Itanium processors showcases HP's commitment to delivering a 64 processor Itanium server and superior investment protection. It is the dawn of a new era in high end computing with the emergence of commodity based hardware.

Superdome supports a multi OS environment. Currently, HP UX, Windows Server 2003, Red Hat, SUSE , and OpenVMS are shipping with Integrity Superdome sx1000. Customers can order any combination of HP UX 11i v2, Windows Server 2003 Datacenter Edition, Windows Server 2003 Enterprise Edition, RHEL AS 3 or 4, SLES 9 or10 or OpenVMS running in separate hard partitions.

With the release of HP Superdome sx2000 with dual-core Itanium 2 processors, HP-UX, Windows Server 2003, OpenVMS and Linux are all supported in separate hard partitions.

The multi OS environment offered by Superdome is listed below.

HP-UX 11i version 2 (sx1000 & sx2000 systems)	• Improved performance over PA-8700 • Investment protection through upgrades from existing Superdomes to next-generation Itanium 2 processors

Overview

Windows Server 2003, Datacenter Edition or Enterprise Edition for Itanium 2
(sx1000 and sx2000 Itanium systems)

- Extension of industry standard-based computing with the Windows operating system further into the enterprise data center
- Increased performance and scalability over 32-bit implementations
- Lower cost of ownership versus proprietary operating system solutions
- Ideal for scale up database opportunities (such as SQL Server 2000 or 2005 (64 bit), Enterprise Edition)
- Ideal for database consolidation opportunities such as consolidation of legacy 32 bit versions of SQL Server 2000 to SQL Server 2005 (64 bit)

NOTE: Windows Server 2003, Enterprise Edition is not supported with mx2 modules

---

Red Hat RHEL AS 3 and 4 and SUSE SLES 9 or 10
(sx1000 Itanium systems only)
Red Hat RHEL AS 4 and SUSE SLES 10
(sx2000 systems)

- Extension of industry standard computing with Linux further into the enterprise data center
- Lower cost of ownership
- Ideal for server consolidation opportunities
- Not supported on Superdome with mx2 processor modules

OpenVMS V8.2-1

- Unmatched availability and disaster tolerance
  - 100% application availability with clustering
  - Up to 800 km between cluster nodes affords maximum Disaster Tolerance
- Impenetrable security
  - Most secure commercial OS
  - Complies with US Federal Government C2 Standards
- Minimum TCO and outstanding investment protection
  - Allowing nodes in a cluster from a workstation to Superdome as well as mixed architecture clusters
  - Server consolidation
  - UNIX portability features allowing easy port to OpenVMS
- HP Superdome sx1000 configurations
  - Minimum version is OpenVMS V8.2-1
  - Mad9M processors only
  - Requires Firmware Version rel_6.0 (or higher). In systems with PA-RISC and Integrity partitions, firmware version rel_6.1 (or higher) is required.
- HP Superdome sx2000 configurations
  - Minimum version is OpenVMS V8.3
  - Dual-core Itanium 2 processors only

Overview

Superdome Service Solutions

Superdome continues to provide the same positive Total Customer Experience via industry-leading HP Services, as with existing Superdome servers. The HP Services component of Superdome is as follows:

- HP customers have consistently achieved higher levels of satisfaction when key components of their IT infrastructures are implemented using the Solution Life Cycle. The Solution Life Cycle focuses on rapid productivity and maximum availability by examining customers' specific needs at each of five distinct phases (plan, design, integrate, install, and manage) and then designing their Superdome solution around those needs. HP offers three pre configured service solutions for Superdome that provides customers with a choice of lifecycle services to address their own individual business requirements.
  - Foundation Service Solution: This solution reduces design problems, speeds time-to-production, and lays the groundwork for long-term system reliability by combining pre-installation preparation and integration services, hands on training and reactive support. This solution includes HP Support Plus 24 to provide an integrated set of 24x7 hardware and software services as well as software updates for selected HP and third party products.
  - Proactive Service Solution: This solution builds on the Foundation Service Solution by enhancing the management phase of the Solution Life Cycle with HP Proactive 24 to complement your internal IT resources with proactive assistance and reactive support. Proactive Service Solution helps reduce design problems, speed time to production, and lay the groundwork for long term system reliability by combining pre installation preparation and integration services with hands on staff training and transition assistance. With HP Proactive 24 included in your solution, you optimize the effectiveness of your IT environment with access to an HP-certified team of experts that can help you identify potential areas of improvement in key IT processes and implement necessary changes to increase availability.
  - Critical Service Solution: Mission Critical environments are maintained by combining proactive and reactive support services to ensure maximum IT availability and performance for companies that can't tolerate downtime without serious business impact. Critical Service Solution encompasses the full spectrum of deliverables across the Solution Lifecycle and is enhanced by HP Critical Service as the core of the management phase. This total solution provides maximum system availability and reduces design problems, speeds time-to-production, and lays the groundwork for long term system reliability by combining pre-installation preparation and integration services, hands on training, transition assistance, remote monitoring, and mission critical support. As part of HP Critical Service, you get the services of a team of HP certified experts that will assist with the transition process, teach your staff how to optimize system performance, and monitor your system closely so potential problems are identified before they can affect availability.
- HP's Mission-Critical Partnership: This service offering provides customers the opportunity to create a custom agreement with Hewlett Packard to achieve the level of service that you need to meet your business requirements. This level of service can help you reduce the business risk of a complex IT infrastructure, by helping you align IT service delivery to your business objectives, enable a high rate of business change, and continuously improve service levels. HP will work with you proactively to eliminate downtime, and improve IT management processes.
- Service Solution Enhancements: HP's full portfolio of services is available to enhance your Superdome Service Solution in order to address your specific business needs. Services focused across multi-operating systems as well as other platforms such as storage and networks can be combined to compliment your total solution.

Minimum/Maximum Configurations for Superdome (sx1000 & sx2000) with Intel Itanium 2 Processors (1.5 GHz and 1.6 GHz)
NOTE: : HP Superdome (sx2000) systems (Itanium 2 (1.6 GHz) "'Mad9M") support only HP UX and Windows.

System Size	Minimum Configuration				Maximum Configuration (One Partition)				
	HP-UX 11i V2	Windows 2003 Datacenter Edition and Enterprise Edition[4]	Red Hat RHEL AS3 & 4and SuSE SLES 9 & 10 (sx1000 only)	OpenVMS version 8.2 1 (or higher) 3	HP-UX 11i V2	Windows 2003 Datacenter/ *Enterprise Edition[4]	SUSE SLES 9 (sx1000 only)	Red Hat RHEL AS 4 (sx1000 only)	OpenVMS version 8.2 1 (or higher)[3]
16 Processors									
Processors	2	2	2	2	16	16/ *8	16	16	16
Memory	2 GB	2 GB	2 GB	2 GB	256 GB/ 512 GB (sx2000)	256 GB (sx1000 with DC Edition)/ 512GB (sx2000 with DC Edition)/ 128 GB (sx1000 with EE)/ 256 GB (sx2000 with EE)	256 GB	256 GB	256 GB
Cell Boards	1 Cell Board	1 Cell Board	1 Cell Board	1 Cell Board	4 Cell Boards	4 Cell Boards/ *2 Cell Boards	4 Cell Boards	4 Cell Boards	4 Cell Boards
PCI-X Chassis	1	1	1	1	4	4/ *2	2	2	4
Max nPARS	Not applicable	Not applicable	Not applicable	Not applicable	4	4/ *2	4	4	4
32-Processors									
Processors	2	2	2	2	32	32/ *8	16	16	16

Standard Features

Memory	2 GB	2 GB	2 GB	2 GB	512 GB/ 256 GB (sx2000)	512 GB (sx1000 with DC Edition)/ 1024 GB (sx2000 with DC Edition)/ 128 GB (sx1000 with EE)/ 256 GB (sx2000 with EE)	256 GB	256 GB	256 GB
Cell Boards	1 Cell Board	1 Cell Board	1 Cell Board	1 Cell Board	8 Cell Boards	8 Cell Boards/ *2 Cell Boards	4 Cell Boards	4 Cell Boards	4 Cell Boards
PCI-X Chassis	1	1	1	1	8	8/ *2	2	2	4
Max nPARS	Not applicable	Not applicable	Not applicable	Not applicable	8 (1)	81/ *2	8(1)	8(1)	8(1)
64 Processors									
Processors	6	6	Not Supported	2	64	64/ *8	Not Supported	Not Supported	16
Memory	6 GB	6 GB		2 GB	1024 GB/ 2048 GB (sx2000)	1024 GB (both sx1000 and sx2000 with DC Edition)/ 128 GB (sx1000 with EE)/ 256 GB (sx2000 with EE)			256 GB
Cell Boards	3 Cell Boards	3 Cell Boards Datacenter Edition; 2 Cell Boards Enterprise Edition		1 Cell Board	16 Cell Boards	16 Cell Boards/ *2 Cell Boards			4 Cell Boards
PCI-X Chassis	1	1	1	1	16	16/ *2	2	2	4
Max nPARS	Not Applicable	Not Applicable	Not Applicable	Not Applicable	16 (2)	162/ *2	16 (2)	16 (2)	16 (2)

[1]Requires I/O expansion cabinet if more than 4 nPARS
[2]Requires I/O expansion cabinet if more than 8 nPARS
[3]1.6 Ghz Mad9M Processors only on sx1000 configurations for OpenVMS
[4]Windows Server 2003, Enterprise Edition is not support with mx2 modules.
*NOTE: HP Superdome (sx2000) systems (Itanium 2 (1.6 GHz) "Mad9M") support only HP UX and Windows.

Minimum/Maximum Configurations for Superdome (sx2000) with Dual-core Intel Itanium 2 Processors

	Minimum Configuration				Maximum Configuration (One Partition)				
	HP UX 11i V2	Windows 2003 Datacenter Edition and Enterprise Edition	Red Hat R4 U3	OpenVMS V8.3 (or higher)[3]	HP UX 11i V2	Windows 2003 Datacenter/ Enterprise Edition	SUSE SLES 10 (Q1 07)	Red Hat RHEL 4 (U3)	OpenVMS version 8.3 (or higher)[3]
16 Processors / 32 Cores									
Processors/ cores	1/2	1/2	1/2	1/2	16/32	Datacenter 16/32; Enterprise 8/16		16/32	16/32
Memory	2 GB	2 GB	4GB	2 GB	512 GB	Datacenter 512 GB; Enterprise256 GB		256 GB	512 GB
Cell Boards	1 Cell Board	1 Cell Board	1 Cell Board	1 Cell Board	4 Cell Boards	Datacenter 4 Cell Boards; Enterprise 2 Cell Boards		4 Cell Boards	4 Cell Boards
PCI X Chassis	1	1	1	1	4	Datacenter 4; Enterprise 2		4	4
Maximum nPars	Not applicable	Not applicable	Not applicable	Not applicable	4	Datacenter 4; Enterprise 2		4	4
32 Processors / 64 Cores									
Processors / Cores	1/2	1/2	1/2	1/2	32/64	Datacenter 32/64; Enterprise 8/16		32/64	16/32
Memory	2 GB	2 GB	4 GB	2 GB	1024 GB	Datacenter 1024 GB; Enterprise 256 GB		512 GB	512 GB
Cell Boards	1 Cell Board	1 Cell Board	1 Cell Board	1 Cell Board	8 Cell Boards	Datacenter 8 Cell Boards; Enterprise 2 Cell Boards		8 Cell Boards	4 Cell Boards
PCI X Chassis	1	1	1	1	8	Datacenter 8; Enterprise 2		4	4

Standard Features

Maximum nPars	Not applicable	Not applicable	Not applicable	Not applicable	8[1]	Datacenter 81; Enterprise 2		8[1]	8[1]
64 Processors / 128 Cores									
Processors / Cores	2/4	2/4	To be supported Q1 07	1/2	64/128	Datacenter 32/64; Enterprise 8/16		To be supported Q1 07	16/32
Memory	6 GB	6 GB		2 GB	2048 GB	Datacenter 1024 GB; Enterprise 256 GB			512 GB
Cell Boards	3 Cell Boards	3 Cell Boards Datacenter Edition; 2 Cell Boards Enterprise Edition		1 Cell Board	16 Cell Boards	Datacenter 16 Cell Boards; Enterprise 2 Cell Boards			4 Cell Boards
PCI X Chassis	1	1		1	16	Datacenter 16; Enterprise 2			4
Maximum nPars	Not applicable	Not applicable	Not applicable	Not applicable	16[2]	Datacenter 162; Enterprise 2			162

[1]Requires I/O expansion cabinet if more than 4 nPars
[2] Requires I/O expansion cabinet if more than 8 nPars
[3] Dual core Montecito processors only on sx2000 configurations for OpenVMS

Standard Hardware Features

Superdome with Intel Itanium 2 and Dual core Itanium 2 Processors Standard Hardware Features

- Redundant Power supplies
- Redundant Fans
- Factory integration of memory and I/O cards
- Installation Guide, Operator's Guide and Architecture Manual
- HP site planning and installation
- One-year warranty with same business day on-site service response

Standard Features

The HP Integrity (sx1000) Superdome servers may require a firmware update to support Intel Itanium 2 Processor/Cell Add on products shipping after June 15, 2005.

Affected Intel Itanium 2 processor products for the Integrity (sx1000) Superdome are:

- AD003A-Intel Itanium 2 1.6 GHz 9 MB processor (2 pack)
- AD004A-Instant Capacity right to access Intel Itanium 2 1.6 GHz 9 MB processor

ACTION:
Check the server firmware prior to installing any of these processor products. The firmware versions can be checked using the MP "sysrev" command. The HP Integrity Superdome requires "Firmware Version rel_5.1" or later. The individual firmware revisions that make up version rel_5.1 and the firmware upgrade instructions are in the Release Notice that is included in the firmware download bundle. If you are not able to check the firmware or if the firmware is not at version rel_5.1 or later, contact HP support.

NOTE: For sx1000 configurations, OpenVMS requires firmware version rel_6.0 or later and for for systems with PA RISC and Integrity partitions, firmware version rel_6.1 or later.

NOTE: Once the firmware is at the supported revision level, proceed with attaching the Processor/Cell Add-On Products to the server using the Service Guide. The Service Guide is available at http://docs.hp.com.

Minimum/Maximum Configurations for Superdome with mx2 Processor Modules (sx1000 systems only)

System Size	Minimum Configuration		Maximum Configuration (One Partition)	
	HP-UX 11i V2	Windows 2003 Datacenter	HP-UX 11i V2	Windows 2003 Datacenter
16 Processors				
Processors	2	2	32	32
Memory	2 GB	2 GB	256 GB	256 GB
Cell Boards	1 Cell Board	1 Cell Board	4 Cell Boards	4 Cell Boards
PCI-X Chassis	1	1	4	4
Max nPARS	Not applicable	Not applicable	4	4
32-Processors				
Processors	2	2	64	64
Memory	2 GB	2 GB	512 GB	512 GB
Cell Boards	1 Cell Board	1 Cell Board	8 Cell Boards	8 Cell Boards
PCI-X Chassis	1	1	8	8
Max nPARS	Not applicable	Not applicable	8 (1)	8 (1)
64-Processors				
Processors	6	6	128	64
Memory	6 GB	6 GB	1024 GB	1024 GB
Cell Boards	3 Cell Boards	3 Cell Boards	16 Cell Boards	16 Cell Boards
PCI-X Chassis	1	1	16	16
Max nPARS	Not Applicable	Not Applicable	16 (2)	16 (2)

Standard Features

Standard Hardware Features	• Redundant Power supplies • Redundant Fans • Factory integration of memory and I/O cards • Installation Guide, Operator's Guide and Architecture Manual • HP site planning and installation • One-year warranty with same business day on-site service response
	(1) Requires I/O Expansion cabinet if more than 4 nPARS (2) Requires I/O Expansion cabinet if more than 8 nPARS

Configuration

There are three basic building blocks in the Superdome system architecture: the cell, the crossbar backplane and the PCI-X based I/O subsystem.

Cabinets (sx1000 & sx2000)

Starting with the sx1000 chip set, Superdome servers will be released with the Graphite color. A Superdome system will consist of up to four different types of cabinet assemblies:

- One Superdome left cabinet.
- No more than one Superdome right cabinet (only Superdome 64-processor system)
  The Superdome cabinets contain all of the processors, memory and core devices of the system. They will also house most (usually all) of the system's PCI-X cards. Systems may include both left and right cabinet assemblies containing, a left or right backplane respectively.
- One or more HP Rack System/E or the new HP Universal 10K G2 cabinets. These 19 inch rack cabinets are used to hold the system peripheral devices such as disk drives.
- Optionally, one or more I/O expansion cabinets (Rack System/E or the new HP Universal 10K G2). An I/O expansion cabinet is required when a customer requires more PCI X cards than can be accommodated in their Superdome cabinets.

Superdome cabinets will be serviced from the front and rear of the cabinet only. This will enable customers to arrange the cabinets of their Superdome system in the traditional row fashion found in most computer rooms. The width of the cabinet will accommodate moving it through common doorways in the U.S. and Europe. The intake air to the main (cell) card cage will be filtered. This filter will be removable for cleaning/replacement while the system is fully operational.

A status display will be located on the outside of the front and rear doors of each cabinet. The customer and field engineers can therefore determine basic status of each cabinet without opening any cabinet doors.

Superdome 16 processor and Superdome 32 processor systems are available in single cabinets. Superdome 64 processor systems are available in dual cabinets.

Each cabinet may contain a specific number of cell boards (consisting of processors and memory) and I/O. See the following sections for configuration rules pertaining to each cabinet.

Cells (Processors and Memory)

A cell, or cell board, is the basic building block of a Superdome system. It is a symmetric multi processor (SMP), containing up to 4 processor modules and up to 16 GB of main memory using 512 MB DIMMs (sx1000 only), up to 32 GB of main memory using 1 GB DIMMs and up to 64 GB of main memory using 2 GB DIMMs. It is also possible to mix 512 MB, 1 GB, and 2 GB DIMMs on the same cell board. A connection to a 12 slot PCI X card cage is optional for each cell.

The Superdome cell boards shipped from the factory are offered with 2 processors or 4 processors. These cell boards are different from those that were used in the previous PA RISC releases of Superdome.

The Superdome cell board contains:

- (sx1000):Itanium 2 1.5 GHz processors or Itanium 2 1.6 GHz processors (up to 4 processor modules for a total of 4 cores) or mx2 dual processor modules (up to 4 modules for a total of 8 cores)
- (sx2000): Dual core Itanium 2 1.6GHz processor modules (up to 4 processor modules for a total of 8 cores) or Itanium 2 1.6 GHz processors (up to 4 processor modules for a total of 4 cores)
- Cell controller ASIC (application specific integrated circuit)
- Main memory DIMMs (up to 32 DIMMs per board in 4 DIMM increments, using 512 MB [sx1000 only], 1 GB, 2 GB, or 4-GB [sx2000 only] DIMMs-or some combination of both. sx2000 based

systems use the 1 GB, 2 GB or 4 GB DIMMs only)
- Voltage Regulator Modules (VRM)
- Data buses
- Optional link to 12 PCI X I/O slots

Crossbar Backplane

Sx2000: Each backplane contains 3 independent crossbar fabrics resulting in a robust highly available system. A High speed serial (HSS) link technology is employed for crossbars interconnect. Each backplane cabinet can support up to eight cells. Eight cells support a 32 processor module Superdome. Two backplanes (two cabinets) can be linked together with flex cables to produce a complex that can support up to 16 cells resulting in a 64 processor module Superdome (64 core Itanium 2 (Mad9M)).
Sx1000: Each crossbar backplane contains two sets of two crossbar chips that provide a non blocking connection between eight cells and the other backplane. Each backplane cabinet can support up to eight cells. A backplane supporting four cells results in a 16 processor Superdome (16 core Itanium 2 Madison or 32 core Itanium2 mx2). Similarly, a backplane supporting eight cells would result in a 32 processor Superdome (32 core Itanium 2 Madison or 64 core Itanium 2 mx2). Two backplanes can be linked together with flex cables to produce a cabinet that can support up to 16 cells resulting in a 64 processor Superdome (128 core Dual core Itanium 2 or 64 core Itanium 2 (Mad9M)).

I/O Subsystem

Each I/O chassis provides twelve I/O slots.
SX1000: Superdome PA 8600 and Superdome PA 8700 support I/O chassis with 12 PCI 66 capable slots, eight PCI 2x, supported via single (1x) ropes (266 MB/s peak) and four PCI 4x, supported via dual (1x) ropes (533 MB/s peak). Superdome PA 8600 and PA 8700 systems support PCI I/O chassis only. Superdome PA 8800/PA 8900 supports I/O chassis with 12 PCI X 133 capable slots, eight PCI 4x, supported via single enhanced (2x) ropes (533 MB/s peak) and four PCI X supported via dual enhanced (2x) ropes (1066 MB/s peak). Superdome PA 8800/PA 8900 will be shipped with PCI X I/O chassis, but supports both PCI and PCI X I/O chassis for in field upgrades.
Sx2000: The sx2000 Superdome supports I/O chassis with two PCI-X 266 MHz slots, six PCI-X 133MHz slots and 4 PCI -X 66 MHz slots. These I/O chassis employ a PCI-X bus for each I/O slot. PCI-X mode 2 is supported.

- Each Superdome cabinet supports a maximum of four I/O chassis. The optional I/O expansion cabinet can support up to an additional six I/O chassis
- A 4 cell Superdome (16 processor) supports up to four I/O chassis for a maximum of 48 PCI X slots.
- An 8 cell Superdome (32 processor) supports up to eight I/O chassis for a maximum of 96 PCI X slots. Four of these I/O chassis will reside in an I/O expansion cabinet.
- A 16 cell Superdome (64 processor) supports up to sixteen I/O chassis for a maximum of 192 PCI X slots. Eight of these I/O chassis will reside in two I/O expansion cabinets (either six chassis in one I/O expansion cabinet and two chassis in the other, or four chassis in each).

Configuration

Core I/O (required by sx1000 systems only, sx2000 systems do not have a core I/O)

The core I/O in Superdome provides the base set of I/O functions required by every Superdome partition. Each partition must have at least one core I/O card in order to boot. Multiple core I/O cards may be present within a partition (one core I/O card is supported per I/O backplane); however, only one may be active at a time. Core I/O will utilize the standard long card PCI-X form factor but will add a second card cage connection to the I/O backplane for additional non-PCI X signals (USB and utilities). This secondary connector will not impede the ability to support standard PCI-X cards in the core slot when a core I/O card is not installed.

Any I/O chassis can support a Core I/O card that is required for each independent partition. A system configured with 16 cells, each with its own I/O chassis and core I/O card could support up to 16 independent partitions. Note that cells can be configured without I/O chassis attached, but I/O chassis cannot be configured in the system unless attached to a cell.

HP-UX Core I/O (A6865A)(sx1000 systems only, sx2000 systems do not have a core I/O)

The core I/O card's primary functions are:

- Partitions (console support) including USB and RS-232 connections
- 10/100Base-T LAN (general purpose)

Other common functions, such as Ultra/Ultra2 SCSI, Fibre Channel, and Gigabit Ethernet, are not included on the core I/O card. These functions are, of course, supported as normal PCI-X add-in cards.

The unified 100Base-T Core LAN driver code searches to verify whether there is a cable connection on an RJ-45 port or on an AUI port. If no cable connection is found on the RJ-45 port, there is a busy wait pause of 150 ms when checking for an AUI connection. By installing the loopback connector (description below) in the RJ-45 port, the driver would think an RJ-45 cable was connected and would not continue to search for an AUI connection, hence eliminate the 150 ms busy wait state

Product/ Option Number	Description
A7108A	RJ-45 Loopback Connector
0D1	Factory integration RJ-45 Loopback Connector

Windows Core I/O (A6865A and optional VGA/USB A6869A/A6869B)- (sx1000 systems only, sx2000 systems do not have core I/O)

For Windows Server 2003, Windows does not support the 10/100 LAN on the A6865A core I/O card, a separate Gigabit Ethernet card such as the A7061A, A7073A, A9899A or A9900A is required. The use of Graphics/USB card (A6869A/A6869B) is optional and not required.
NOTE: A6869A is not supported on systems with the sx2000 chipset.

Linux Core I/O (A6865A)(sx1000 systems only, sx2000 systems do not have core I/O)

The core I/O card's primary functions are:

- Partitions (console support) including USB and RS-232 connections
- 10/100Base-T LAN (general purpose)

Other common functions, such as Ultra/Ultra2 SCSI, Fibre Channel, and Gigabit Ethernet, are not included on the core I/O card. These functions are supported as normal PCI-X add-in cards.

Configuration

OpenVMS Core I/O (A6865A)(sx1000 systems only, sx2000 systems do not have core I/O)

The core I/O card's primary functions are:

- Partitions (console support) including USB and RS 232 connections
- 10/100Base T LAN (general purpose)

Other common functions, such as Ultra320 SCSI, Fibre Channel, and Gigabit Ethernet, are not included on the core I/O card. These functions are supported as normal PCI X add in cards.

I/O Expansion Cabinet

The I/O expansion functionality is physically partitioned into four rack-mounted chassis—the I/O expansion utilities chassis (XUC), the I/O expansion rear display module (RDM), the I/O expansion power chassis (XPC) and the I/O chassis enclosure (ICE). Each ICE supports up to two 12-slot PCI-X chassis.

Field Racking

The only field rackable I/O expansion components are the ICE and the 12-slot I/O chassis. Either component would be field installed when the customer has ordered additional I/O capability for a previously installed I/O expansion cabinet.

No I/O expansion cabinet components will be delivered to be field installed in a customer's existing rack other than a previously installed I/O expansion cabinet. The I/O expansion components were not designed to be installed in racks other than Rack System E or the new Universal 10K G2 rack. In other words, they are not designed for Rosebowl I, pre merger Compaq, Rittal, or other third party racks.

The I/O expansion cabinet is based on a modified HP Rack System E and the new Universal 10K G2 and all expansion components mount in the rack. Each component is designed to install independently in the rack. The Rack System E and the Universal 10K G2 cabinet have been modified to allow I/O interface cables to route between the ICE and cell boards in the Superdome cabinet. I/O expansion components are not designed for installation behind a rack front door. The components are designed for use with the standard Rack System E and the Universal 10K G2 perforated rear door.

I/O Chassis Enclosure (ICE)

The I/O chassis enclosure (ICE) provides expanded I/O capability for Superdome. Each ICE supports up to 24 PCI-X slots by using two 12-slot Superdome I/O chassis. The I/O chassis installation in the ICE puts the PCI-X cards in a horizontal position. An ICE supports one or two 12-slot I/O chassis. The I/O chassis enclosure (ICE) is designed to mount in a Rack System E rack and consumes 9U of vertical rack space.

To provide online addition/replacement/deletion access to PCI or PCI-X cards and hot-swap access for I/O fans, all I/O chassis are mounted on a sliding shelf inside the ICE.

Four (N+1) I/O fans mounted in the rear of the ICE provide cooling for the chassis. Air is pulled through the front as well as the I/O chassis lid (on the side of the ICE) and exhausted out the rear. The I/O fan assembly is hot swappable. An LED on each I/O fan assembly indicates that the fan is operating.

Cabinet Height and Configuration Limitations

Although the individual I/O expansion cabinet components are designed for installation in any Rack System E cabinet, rack size limitations have been agreed upon. IOX Cabinets will ship in either the 1.6 meter (33U) , 1.96 meter (41U) or the 2.0-meter (Universal 10K G2)(42U) cabinet. In order to allay service access concerns, the factory will not install IOX components higher than 1.6 meters from the floor. Open space in an IOX cabinet will be available for peripheral installation.

Configuration

Peripheral Support

All peripherals qualified for use with Superdome and/or for use in a Rack System E or Universal 10K G2 are supported in the I/O expansion cabinet as long as there is available space. Peripherals not connected to or associated with the Superdome system to which the I/O expansion cabinet is attached may be installed in the I/O expansion cabinet.

Server Support

No servers except those required for Superdome system management such as Superdome Support Management Station or ISEE may be installed in an I/O expansion.

Peripherals installed in the I/O expansion cabinet cannot be powered by the XPC. Provisions for peripheral AC power must be provided by a PDU or other means.

Standalone I/O Expansion Cabinet

If an I/O expansion cabinet is ordered alone, its field installation can be ordered via option 750 in the ordering guide (option 950 for Platinum Channel partners).

DVD Solution

The DVD solution for Superdome requires the following components. These components are required per partition. External racks A4901A and A4902A must also be ordered with the DVD solution. Note: One DVD is required and one DAT is recommended per nPartition.

NOTE: One DVD and one DAT is required per nPartition.

NOTE: For OpenVMS, only the one DVD is required per nPartition.

Superdome DVD Solutions

Description	Part Number	Option Number
PCI 2 channel Ultra320 SCSI (required for OpenVMS)	A7173A	0D1
Surestore Tape Array 5300	C7508AZ	
HP DVD+RW Array Module (one per partition) NOTE: The HP DVD-ROM Array Module for the TA5300 (C7499B) is replaced by HP DVD+RW Array Module (Q1592A) to provide customers with read capabilities for loading software from CD or DVD, DVD write capabilities for small amounts of data (up to 4 GB) and offline hot-swap capabilities. Windows supports using and reading from this device, but Windows does not support DVD write with this device. OpenVMS supports reading and booting, only, from this device, at this time.	Q1592A	0D1
DDS-4/DAT40 (DDS-5/DAT 72 is also supported. Product number is Q1524B) (one per partition)	C7497B	0D1
Jumper SCSI Cable for DDS-4 (optional)[1]	C2978B	0D1
SCSI cable 1-meter multi-mode VH-HD68	C2361B	0D1
SCSI Terminator	C2364A	0D1

[1]0.5-meter HD HDTS68 is required if DDS-4 or DDS-5 is used.
[2]Supported, but may no longer be orderable.

Configuration

Partitions

Superdome can be configured with hardware partitions, (nPars). Given that Windows Server 2003, SuSE SLES 9, Red Hat (RHEL) AS 3 and 4, and OpenVMS do not support virtual partitions (vPars), Superdome systems running Windows Server 2003, Datacenter Edition or *Enterprise Edition, SuSE SLES 9, Red Hat RHEL AS 3 and 4, or OpenVMS, do not support vPars.

A hardware partition (nPar) consists of one or more cells that communicate coherently over a high bandwidth, low latency crossbar fabric. Individual processors on a single cell board cannot be separately partitioned. Hardware partitions are logically isolated from each other such that transactions in one partition are not visible to the other hardware partitions within the same complex.

Each nPar runs its own independent operating system. Different nPars may be executing the same or different revisions of an operating system, or they may be executing different operating systems altogether. Superdome supports HP UX 11i version 2, Windows Server 2003, Datacenter Edition or Enterprise Edition, SuSE SLES 10, Red Hat RHEL AS 3 and 4, and OpenVMS operating systems. The diagram below shows a multi OS environment within Superdome. Note: The sx2000 Superdome with Mad9M processors only supports HP UX 11i version2 and Windows Server 2003, Datacenter Edition or Windows Server 2003, Enterprise Edition.

NOTE: It is possible for PA 8800 and PA 8900 nPars to co exist with Itanium 2 1.6 GHz nPars in the same Superdome system, but on different partitions. Customers can configure an Itanium 2 1.6 GHz nPar in an HP 9000 Superdome running PA 8800 or PA 8900 (and vice versa) in the field only. Factory orders for mixed Itanium and PA RISC nPars are not allowed.
*Enterprise Edition is not supported with mx2 modules.

Each nPar has its own independent processors, memory and I/O resources consisting of the resources of the cells that make up the partition. Resources (cell boards and/or I/O chassis) may be removed from one nPar and added to another without having to physically manipulate the hardware, but rather by using commands that are part of the System Management interface. The table below shows the maximum size of nPars per operating system:

	HP-UX 11i Version 2	Windows Server 2003	Red Hat	SUSE SLES 9	OpenVMS
Maximum size of nPar	64 processors/128 cores, 512 GB RAM	32 processors/64 cores, Data Center 8 processors/16 cores Enterprise Edition 512 GB RAM	8 processors (RHEL 3), 32 processors, 64 cores (RHEL4) 512 GB RAM	16 processors,32 cores 256 GB RAM	sx1000 - 4 cells with up to 16 P/16C, 256 GB RAM sx2000 - 4 cells with up to 16 Processors/ 32 Cores, 512 GB RAM
Maximum number of nPars	16	16	16	16	16

For information on type of I/O cards for networking and mass storage for each operating environment, please refer to the **Technical Specifications** section of this document. For licensing information for each operating system, please refer to the Ordering Guide.

Superdome supports static partitions. Static partitions imply that any nPar configuration change requires a reboot of the nPar. In a future HP-UX and Windows release, dynamic nPars will be supported. Dynamic npars imply that nPar configuration changes do not require a reboot of the nPar. Using the related capabilities of dynamic reconfiguration (i.e. on-line addition, on-line removal), new resources may be added to an nPar and failed modules may be removed and replaced while the nPar continues in operation. Adding new nPars to Superdome system does not require a reboot of the system.

Windows Server 2003, Datacenter edition for Itanium-based systems - HP Product Structure

Product Number T2372A
Pre-loaded Windows Server 2003, Datacenter Edition for Itanium 2 systems

Options:

- 0D1 - factory integration
- B01 - on site installation at customer's location (must contact HP Services for a quote to install on-site!)
- ABA - English localization only (other languages, Ger, Fre, Ita available only as a special with extra lead time)
- ABJ - Japanese localization
- 002 - 2 processor LTU
- 004 - 4 processor LTU
- 008 - 8 processor LTU
- 016 - 16 processor LTU
- 032 - 32 processor LTU
- 064 - 64 processor LTU

Windows Server 2003, Enterprise Edition for Itanium based Systems- HP Product Structure

Product Number T2373A
Pre-loaded Windows Server 2003, Enterprise Edition for Itanium 2 systems

Options:

- 0D1 - factory integration
- ABA - English localization only (other languages, Ger, Fre, Ita available only as a special with extra lead time)

NOTE: See Windows 2003 Ordering Section for further information. Not available/supported in Japan. Windows Server 2003, Enterprise Edition is not supported with mx2 modules.

Mixing of PA-RISC and Itanium Cells in Superdome

1) Which processors can be mixed in a Superdome?

The first step in determining which processors can be mixed within a Superdome is to look at the chipset. A Superdome can only support one type of chipset (legacy chipset, sx1000 chipset or sx2000 chipset) at a time.

The legacy chipset only supports the PA-RISC architecture. The PA-8600, PA-8700, and PA-8700+ processors were supported with this chipset. As a result, they can be mixed within the Superdome but they can not be mixed with processors supported by other chipsets (i.e. Itanium 2 9M with the sx1000 chipset).

With the sx1000 chipset, processors of like architectures (PA-RISC and Itanium 2 architectures) that are supported by the sx1000 chipset can be mixed in separate hard partitions. For example, the HP 9000 Superdome supports mixing the PA-8800 and PA-8900 processors in separate hard partitions. The HP Integrity Superdome supports mixing the Itanium 2 6M, and Itanium 2 9M processors, as well as the mx2 dual processor module in separate hard partitions. In addition, a subset of the PA-RISC and Itanium 2 processors (PA-8800, PA-8900 and Itanium 2 9M processors) can be supported at the same time in different hard partitions within a Superdome.

With the sx2000 chipset, Dual core Itanium 2 and Itanium2 1.6 MHz processors are the only supported processors and therefore you cannot mix other processors in this system.

The table below highlights which processors can co-exist on a Superdome in separate hard partitions.

	PA 8600	PA 8700	PA 8700 +	PA 8800	PA 8900	Itanium 2 6M	Itanium 2 9M	mx2 dual processor module	Dual core Itanium 2
PA 8600	NA	Yes	Yes	No	No	No	No	No	No
PA 8700	Yes	NA	Yes	No	No	No	No	No	No
PA 8700 +	Yes	Yes	NA	No	No	No	No	No	No
PA 8800	No	No	No	NA	Yes	No	Yes	No	No
PA 8900	No	No	No	Yes	NA	No	Yes	No	Yes
Itanium 2 6M (sx1000)	No	No	No	No	No	NA	Yes	Yes	No
Itanium 2 9M (sx1000)	No	No	No	Yes	Yes	Yes	NA	Yes	No
mx2 dual processor module	No	No	No	No	No	Yes	Yes	NA	No
Itanium 2 9M (sx2000)	No	No	No	No	No	No	N/A	No	Yes
Dual core Itanium 2 (sx2000)	No	No	No	No	Yes (future)	No	Yes	No	NA

2) In order to run an Itanium-based partition in an Integrity server, what changes are required?

In order to add a new partition with Itanium 2 9M processors on an HP 9000 Superdome, the following steps are required:

Step #1: Upgrade firmware on PA-RISC based partitions

Step #2: Create a new hard partition in the Superdome for Itanium 2-based cell (s)

Step #3: Plug in cell boards for Itanium 2-based cells

Step #4: Some I/O cards may need to be added for that specific hard partition (Windows does not support the identical set of I/O cards that HP-UX 11i supports)

Step #5: Load operating system for Itanium 2-based partition

Upgrading a PA-RISC partition to support Itanium 2 processors would require similar steps:

Step #1: Upgrade firmware on PA-RISC based partitions

Step #2: Pull out existing PA-RISC cell boards

Step #3: Swap existing memory into cell boards for Intel Itanium 2 processor (protects investment in current memory)

Step #4: Plug in cell boards for Itanium 2-based cell boards

Step #5: Some I/O cards may need to be added for that specific hard partition (Windows does not support the identical set of I/O cards that HP-UX 11i supports)

Step #6: Load operating system for Itanium 2-based partition

The in-box addition of Itanium 2 processors can be done with no additional hardware, no new chassis and no change to backplane.

3) When will mixing of PA-RISC and Itanium 2 processors be available?

Support for mixing PA RISC and Itanium 2 processors with the sx1000 chipset in separate hard partitions are available.
NOTE: At this time, OpenVMS on Integrity is not supported on a Superdome that includes PA RISC nPars.
NOTE: At this time, sx2000 chipset does not support PA-RISC processors.

4) Is mixing of PA-RISC and Itanium processors factory configurable?

NO. The plan is to enable PA-RISC systems that shipped from the factory to add Itanium partitions later on in the field. In addition, HP offers the ability to enable Itanium-based systems that shipped from the factory to add PA-RISC partitions later on in the field (this will happen less frequently)

The plan is not to allow PA RISC or Intel Itanium based partitions to be added to systems that shipped from the factory. The sx2000 Superdome does not support PA-RISC processors yet so therefore mixing of Itanium 2 and PA-RISC processors is not supported in the sx2000 systems. Mixing of sx1000 and sx2000 chipsets in the same system is NOT supported.

5) How long does it take to add an Itanium-based partition in an HP 9000 customers?

The effort required to add an Itanium-based partition to an HP 9000 Superdome is slightly more than what is required to add a new PA-RISC-based partition. In both cases you have to create a new hard partition in the Superdome, plug in cell boards for the new processors and load the operating system. There are no hardware changes required to support mixing of processor types. The only known difference at this time is that a customer would need to upgrade the firmware to support the Itanium architecture. The process is currently being tested by R&D. This document will be updated as we learn more through testing. The sx2000 chipset is NOT able to be installed in a sx1000 system. The chipsets in a system must be the same.

6) Will mixing impact the overall performance of a Superdome (i.e. would an Itanium-based partition have the same performance in a mixed Superdome as in a Superdome only populated with Itanium 2 processors)?

No. For example, suppose you have a hard partition with 16 Intel Itanium 2 9M processors. This partition would have the same performance in a system with a mix of PA-RISC and Itanium 2-based hard partitions as a server where all of the other partitions are Itanium 2-based. This is assuming that besides the type of processors, all of the other configuration options are the same (memory, cell board locations, number of partitions, etc).

7) Can I add PA RISC processors to an Itanium based Integrity Superdome?(sx1000 based sytems only, sx2000 systems does not allow mixing at this point)

Yes. HP offers the ability to enable Itanium-based systems that shipped from the factory to add PA-RISC partitions later on in the field. This serves as a safety net/insurance policy. If the customer determines after the fact that a specific application is not ready for the Itanium technology, they can simply add a PA-RISC

hard partition to that Integrity Superdome and continue the execution of the application while it is certified.

8) Why can't PA-8700+ processors be mixed with Itanium 2 9M processors?

For years, HP has had a high level of investment protection with the HP 9000 Superdome. Customers have been provided with the maximum investment protection level by keeping their investment in PA8600 and PA8700 processors while adding the higher speed PA8700+. In addition, each processor type runs at its rated speed so there is no downgrading of the higher speed processors (some competitors are erroneously saying Superdome has to downgrade the higher speed processors). For example, the PA8700 runs at a full 750 MHz while the PA8600 runs at 550 MHz-there is no forcing of the PA8700 down to 550 MHz just to obtain compatibility. This is possible because the processor speeds are independent from the bus speeds.

The Intel Itanium 2 processor family is an advanced architecture featuring exceptional floating point and SSL performance. The Itanium 2 processors have the processing power that is consistent with a Superdome class high-end server. In addition, the Itanium processor accesses cache memory using a data block that is greater than twice the size of the PA8600, PA8700, and PA8700+. Thus, the Itanium processor family and subsequent PA-8800 and PA-8900 processors require a different chipset than the PA8600, PA8700, or PA8700+ processors. One of the key restrictions in mixing processors is that all cell boards in a given Superdome must have the same chipset. The PA-8700+ is supported by a legacy chipset that only supported PA-RISC processors. As a result, the PA-8700+ can only be mixed with PA-8700 and PA-8600 processors. In order to support Itanium 2 9M processors in their current HP 9000 Superdome, the customer would be required first to upgrade all of their PA-8700+ processors to PA-8800 or PA-8900 processors utilizing the sx1000 chipset. Once this is complete, the customer would then have the option of adding the Itanium 2 9M processor in a separate partition because it is also supported by the sx1000 chipset.

9) Why are Itanium 2 6M processors and mx2 dual processor modules not supported?

The original plan was to enable PA-RISC systems that shipped from the factory to add Itanium-based partitions later on. The plan was not to allow PA-RISC-based partitions to be added to systems that shipped from the factory as Itanium-based servers with Itanium 2 6M processors and/or mx2 dual processor modules. This is the only case where we encounter a problem with no support for Itanium 2 6M processors and mx2 dual processor modules. If a customer has already made the decision to go with Integrity Superdomes, it is very unlikely that they will take a back step to PA-RISC-based partitions. Conversely, if a customer wants to upgrade a PA-RISC-based partition to be Itanium 2-based, they are much more likely to use the Itanium 2 9M processor than the Itanium 2 6M processors or mx2 dual processor modules because of the added performance and larger cache sizes.

Finally, this feature requires a significant amount of resources to test the different type of configurations supported in a Superdome. In order to provide this functionality in a timely basis to customers, we had to narrow the scope. Thus, the decision was made to not support Itanium 2 6M processors and mx2 dual processors in a mixed processor type of system.
If a customer with Itanium 2 6M processors or mx2 dual processor modules wants to add PA-RISC modules, they would need to first upgrade the partitions with Itanium 2 6M processors or mx2 dual processor modules to Itanium 2 9M processors and then they could add a PA-RISC-based partition to the Superdome.

10) Which versions of the HP-UX 11i operating system will be supported?

Currently HP plans to support HP-UX 11.11 (HP-UX 11i v1) for the PA-RISC architecture as well as 11.23 (HP-UX 11i v2) for Itanium processors and
PA-RISC processors in a mixed environment. Based on HP's testing, HP believes that using HP-UX 11.11

(HP-UX 11i v1) in a mixed configuration is going to require a PA-RISC Firmware update.

NOTE: HP-UX 11iV2 is the only HP-UX version supported on sx2000 based systems.

11) Does a customer have to power down when they add an Itanium-based partition to a PA-RISC-based Superdome?

It depends.

If the existing PA-RISC based partitions are running HP-UX 11i v2 and PDC 22.1 (released in December 2004) then an Itanium-based partition may be added while the PA-RISC partitions are active. If the customer is running WLM it must be patched for proper operation in a mixed environment. We are currently checking with the WLM team to see what impact this would have on the PA-RISC partitions.

HP's shipping firmware for PA-RISC processors does not fully support mixing if the PA-RISC partitions are running HP-UX 11i v1. In this case, downtime for all PA-RISC based partitions would be required to update to the new PDC. The partitions could all be updated independently if desired. HP will release this new PDC when testing is complete for mixing processor types. The above comment regarding WLM also applies.

12) What are the minimum firmware requirements for mixing various operating environments?

The matrix below describes the minimum firmware requirements for sx1000 systems:

Mixing Mad9M with PA 8800/PA 8900	• 11i v2 HWE0409; 11i v2 HWE0505 (no vPars) • Windows Server 2003 • Linux	• 11i v2 HWE0505 (vPars)	• OpenVMS
11i v1	• PDC 22.3 • IPF 2.54 • MFW 15.14	• PDC 22.3 • IPF 3.x • MFW 15.20	• PDC 22.3 • IPF 3.88, MFW 15.22
11i v2	• PDC 22.1 • IPF 2.50 • MFW 15.14	• PDC 22.1 • IPF 3.x • MFW 15.20	• PDC 22.1 • IPF 3.88, MFW 15.22

Mixing Scenarios: operating systems,

The following table summarizes the various operating environments supported on the PA-8800, PA-8900 and Intel Itanium 2 1.6 GHz processors:

Configuration

firmware and management tools (sx1000 systems only)

Processor	Operating System
PA8800	HPUX 11.i v1 HWE 0312 (with patches) HPUX 11.i v1 HWE 0406 HPUX 11.i v1 HWE 0412 HPUX 11.i v2 HWE 0409 HPUX 11.i v2 HWE 0505
PA8900	HPUX 11.i v1 HWE 0412 HPUX 11.i v2 HWE 0409 HPUX 11.i v2 HWE 0505
Itanium 2 1.6 GHz	HPUX 11.i v2 HWE 0409 HPUX 11.i v2 HWE 0505 Linux RHEL3 U3 (and higher) Linux SLES 9 (and higher) Windows 2003 Server, Smart Setup 3.1 (and higher) OpenVMS V8.2-1 (and higher). [See note below]

NOTE: For 8.2-1, an OpenVMS patch kit, that benefits performance for OpenVMS on all cell based systems, should be applied. The patch kit, VMS821I_IPMI-V0100.ZIPEXE (or any patch kit that supercedes this patch kit), can be downloaded at http://www.itrc.hp.com/service/patch/mainPage.do.

The following table lists the minimum firmware requirements for mixing:

Processor	Supported	Not Supported
PA8800	PDC 22.1 (SMS rel_5.0) or above	PDC 20.8, 21.2
PA8900	PDC 22.1 (SMS rel_5.0) or above	n/a
Itanium 2 1.6 GHz A1	IPF 2.50 (SMS rel_5.0) or above	n/a
Itanium 2 1.6 GHz A2	IPF 2.52 (SMS rel_5.1) or above	n/a

NOTE: PDC 20.8 and PDC 21.2 do not contain the required enhancements for mixing. Customers who are running on either of these PDC versions will be required to upgrade to PDC 22.1 to enable mixing. PDC 22.1 also happens the minimum firmware required to support PA8900.

Any of the above operating environments may be used on a mixed Superdome. In the process of testing all of the above combinations some issues and restrictions were discovered with WLM, gWLM and the GUI version of partition manager on HPUX 11i v1. Firmware, operating system or application updates can resolve these issues. However, such updates may not always be possible or practical.

The following mixing scenarios serve to illustrate the issues that were found during testing and possible ways to work around or resolve them. For simplicity, the first three scenarios assume that all the PA partitions in the Superdome are running the same version of HPUX. The fourth scenario explains the more general case where the PA partitions are running different versions of HPUX.

Scenario 1: PA8800 Partitions running HPUX 11i v1 HWE 0312 or 0406 mixed with Itanium 2 1.6 GHz

Known issue with the GUI version of Parmgr: Both of these HWEs shipped with Parmgr V1. This version of Parmgr and the stack that sits under it are not capable of managing Integrity partitions. Parmgr generates error messages when trying to access the Integrity cells and display information about them.

Alternatives to running Parmgr GUI on the PA partitions:
Integrity partitions (and PA partitions) may be managed from the SMS
Integrity partitions (and PA partitions) may be managed by parcmds from any partition
Integrity partitions (and PA partitions) may be managed by the GUI version of Parmgr on any Integrity partition
The minimum firmware listed in the table above is sufficient for running this configuration.

Possible upgrades: If it is important to be able to use the Parmgr GUI on the PA HPUX 11i v1 you can upgrade HPUX to HWE 0412 and follow the process in Scenario 2. If a customer upgrades to any 11i v1 HWE beyond HWE 0412 (HWE 0509 and later) then the proper nPar provider is already included in the HWE and the firmware upgrades in Scenario 2 are sufficient to get full functionality.

Scenario 2: PA-8800 partitions or PA-8900 partitions running 11i v1 HWE 0412 mixed with Itanium 2 1.6 GHz

Known issue with the GUI version of Parmgr: HPUX 11i v1 HWE 0412 introduced Parmgr V2. This version of Parmgr will exhibit the same issues as Scenario 1 with the minimum required firmware from the table. However, upgrades to firmware and the nPar provider will make this version of Parmgr fully functional and capable of managing Integrity partitions. See the upgrade section below for details.

Alternatives to running Parmgr GUI on the PA partitions: These are essentially the same as Scenario 1.
Integrity partitions (and PA partitions) may be managed from the SMS
Integrity partitions (and PA partitions) may be managed by parcmds from any partition
Integrity partitions (and PA partitions) may be managed by the GUI version of Parmgr on any Integrity partition
The minimum firmware listed in the table above is sufficient for running this configuration.

Possible upgrades: If it is important to be able to use the Parmgr GUI on the PA partitions to manage the Integrity partitions the following updates must be made to the system:
PDC must be at revision 22.3 or higher in the 11.i v1 partition. This version of PDC contains a single fix on top of PDC 22.1 that compensates for the byte reversal in the cell info structure between PA and Integrity (this is due to the different "endianness" of the processors). PDC 22.3 will be released as part of SMS rel_6.0 in September 2005.
Integrity firmware must be upgraded to the version that supports vPars. This version of Integrity firmware contains a fix that populates I/O slot information in cell info structures used by the PA partition management stack. Currently shipping Integrity firmware does not populate this information and causes the PA management stack to report errors for Integrity cells. This version of Integrity firmware will be released as part of SMS rel_6.0 in September 2005.
The nPartition provider ("NPar" bundle) must be updated to version B.11.11.01.04 or above. This version will first be released on HPUX 11.i v1 HWE 0509. The version of nPar provider that shipped with HWE 0412 did not properly handle the single core Madison 9M processors.

Scenario 3: PA-8800 partition or PA-8900 partition running 11.i v2 (HWE 0409 or 0505) mixed with Itanium 2 1.6 GHz

This is the simplest scenario. The minimum firmware listed in the table is sufficient and all the partition management tools just work. The PA partition can manage other PA partitions and the Integrity partition and vice versa. Enjoy.

Scenario 4: PA-8800 or PA-8900 partitions running a mixture of the above scenarios mixed with Itanium 2 1.6 GHz

The previous scenarios apply to the individual partition pairs. For instance, say you have a Superdome with a PA8800 11.i v1 HWE0406 partition, a PA8900 11.i v1 HWE0412 partition a PA8800 11.i v2 partition and an Integrity partition. As described in Scenario 1 the PA8800 11.i v1 HWE0406 partition will not be able to recognize the Integrity cells. Manage the Integrity partition from the SMS, the PA8800 11.i v2 partition or using parcmds. As described in Scenario 2, the PA8900 11.i v1 HWE0412 partition cannot manage the Integrity partition with out the firmware and nPar provider upgrades outlined above. Use the SMS, the PA 11.i v2 partition or parcmds to manage the Integrity partition.

WLM:

WLM (workload manager) currently uses the Unix command "uname -i" to ensure that all the partitions it is managing are on the same Superdome. WLM verifies that all the managed partitions "uname -i" values match. It turns out the Integrity partitions and PA partitions on the same Superdome do not return the same value for "uname -i". This prevents WLM from properly managing all the partitions on a single Superdome in a mixed environment with iCAP. The WLM team has provided patches WLM A.03.00 for both 11.i v1 (PHSS_33499 - s700_800 11.11 WLM A.03.00 Cumulative Patch) and 11.i v2 (PHSS_33477 - s700_800 11.23 WLM A.03.00 Cumulative Patch) to resolve this issue. Customers who wish to use WLM in a mixed environment with iCAP will need to install/upgrade to version A.03.00 and install the appropriate patch. Customers who are using either WLM A.02.x or A.03.x in a non-iCAP environment can use their current version of WLM with no upgrade/patch. Future versions of WLM (A.03.00.01 and later) will not require a patch.

gWLM:

gWLM 1.1.1 (global workload manager) was first released with HPUX 11.i v2 HWE0505. It has the same issue with "uname -i" as WLM. This only affects gWLM iCAP SRDs. VPAR, PSET and FSS SRDs work correctly in mixed environments. The gWLM team plans to remedy this issue in the first maintenance release of gWLM. Contact gwlmfeedback@rsn.hp.com for a workaround if you have a gWLM 1.1.1 customer who needs to deploy an iCAP SRD on a mixed complex before the first maintenance release is available.

Global Workload Manager (gWLM V1.1) is also available for OpenVMS on Integrity Servers.

Single System Reliability/Availability Features

Superdome high availability offering is as follows:

NOTE: Online addition/replacement for cell boards is not currently supported and will be available in a future HP UX release. Online addition/replacement of individual processors and memory DIMMs will never be supported.)

- Processor: The features below nearly eliminate the down time associated with processor cache errors (which are the majority of processor errors). If a processor is exhibiting excessive cache errors, HP UX 11i version 2 will ONLINE activate to take its place. Furthermore, the processor cache will automatically be repaired on reboot, eliminating the need for a service call. Dynamic processor resilience with Instant Capacity enhancement.
  NOTE: Dynamic processor resilience and Instant Capacity are not supported when running Windows Server 2003, SUSE SLES 9, Red Hat, or OpenVMS V8.2 1 in the partition. Cell iCAP support is available with OpenVMS partitions starting with OpenVMS V8.3.
  - Processor cache ECC protection and automatic de allocation
  - Processor bus parity protection
  - Redundant DC conversion
- Memory: The memory subsystem design is such that a single SDRAM chip does not contribute more than 1 bit to each ECC word. Therefore, the only way to get a multiple-bit memory error from SDRAMs is if more than one SDRAM failed at the same time (rare event). The system is also resilient to any cosmic ray or alpha particle strike because these failure modes can only affect multiple bits in a single SDRAM. If a location in memory is "bad", the physical page is deallocated dynamically and is replaced with a new page without any OS or application interruption. In addition, a combination of hardware and software scrubbing is used for memory. The software scrubber reads/writes all memory locations periodically. However, it does not have access to "locked down" pages. Therefore, a hardware memory scrubber is provided for full coverage. Finally data is protected by providing address/control parity protection.
  - Memory DRAM fault tolerance, i.e. recovery of a single SDRAM failure
  - DIMM address / control parity protection
  - Dynamic memory resilience, i.e. page de allocation of bad memory pages during

operation.
NOTE: Dynamic memory resilience is not supported when running Windows Server 2003, SUSE SLES 9,r Red Hat RHEL AS 3, or OpenVMS in the partition.
  - Hardware and software memory scrubbing
  - Redundant DC conversion
  - Cell ICAP.
    NOTE: Cell ICAP is not supported when Windows Server 2003 SUSE SLES 9, Red Hat RHEL AS 3, or OpenVMS V8.2-1 is running in the partition. Cell iCAP support is available with OpenVMS partitions starting with OpenVMS V8.3.

- I/O: Partitions configured with dual path I/O can be configured to have no shared components between them, thus preventing I/O cards from creating faults on other I/O paths. I/O cards in hardware partitions (nPars) are fully isolated from I/O cards in other hard partitions. It is not possible for an I/O failure to propagate across hard partitions. It is possible to dynamically repair and add I/O cards to an existing running partition.
  - Full single-wire error detection and correction on I/O links
  - I/O cards fully isolated from each other
  - HW for the Prevention of silent corruption of data going to I/O
  - On-line addition/replacement (OLAR) for individual I/O cards, some external peripherals, SUB/HUB.
    NOTE: Online addition/replacement (OLAR) is not supported when running Red Hat RHEL AS 3,SUSE SLES 9, or OpenVMS in the partition.
  - Parity protected I/O paths
  - Dual path I/O
- Crossbar and Cabinet Infrastructure:
  - Recovery of a single crossbar wire failure
  - Localization of crossbar failures to the partitions using the link
  - Automatic de-allocation of bad crossbar link upon boot
  - Redundant and hotswap DC converters for the crossbar backplane
  - ASIC full burn-in and "high quality" production process
  - Full "test to failure" and accelerated life testing on all critical assemblies
  - Strong emphasis on quality for multiple-nPartition single points of failure (SPOFs)
  - System resilience to Management Processor (MP)
  - Isolation of nPartition failure
  - Protection of nPartitions against spurious interrupts or memory corruption
  - Hot swap redundant fans (main and I/O) and power supplies (main and backplane power bricks)
  - Dual power source
  - Phone-Home capability
- "HA Cluster-In-A-Box" Configuration: The "HA Cluster-In-A-Box" allows for failover of users' applications between hardware partitions (nPars) on a single Superdome system. All providers of mission critical solutions agree that failover between clustered systems provides the safest availability-no single points of failures (SPOFs) and no ability to propagate failures between systems. However, HP supports the configuration of HA cluster software in a single system to allow the highest possible availability for those users that need the benefits of a non-clustered solution, such as scalability and manageability. Superdome with this configuration will provide the greatest single system availability configurable. Since no single system solution in the industry provides protection against a SPOF, users that still need this kind of safety and HP's highest availability should use HA cluster software in a multiple system HA configuration. Multiple HA software clusters can be configured within a single Superdome system (i.e., two 4-node clusters configured within a 32-processor Superdome system).
  - HP-UX: Serviceguard and Serviceguard Extension for RAC
  - Windows Server 2003: Microsoft Cluster Service (MSCS) - limited configurations supported
  - RHEL 4 and SUSE SLES 9: Serviceguard for Linux
  - HP OpenVMS Cluster Software

Multi-system High Availability

HP-UX 11i v2

Any Superdome partition that is protected by Serviceguard or Serviceguard Extension for RAC can be configured in a cluster with:

- Another Superdome with like processors (i.e., both Superdomes must have Itanium 2 1.6 GHz processors or both Superdomes must have mx2 processor (sx1000 systesm only) modules in the partitions that are to be clustered together.)
- One or more standalone non Superdome systems with like processors
- Another partition within the same single cabinet Superdome (refer to "HA Cluster-in-a-Box" above for specific requirements) that has like processors

Separate partitions within the same Superdome system can be configured as part of different Serviceguard clusters.

Geographically Dispersed Cluster Configurations

The following Geographically Dispersed Cluster solutions fully support cluster configurations using Superdome systems. The existing configuration requirements for non-Superdome systems also apply to configurations that include Superdome systems. An additional recommendation, when possible, is to configure the nodes of cluster in each datacenter within multiple cabinets to allow for local failover in the case of a single cabinet failure. Local failover is always preferred over a remote failover to the other datacenter. The importance of this recommendation increases as the geographic distance between datacenters increases.

- Extended Campus Clusters (using Serviceguard with MirrorDisk/UX)
- Metrocluster with Continuous Access XP
- Metrocluster with EMC SRDF
- Continentalclusters

From an HA perspective, it is always better to have the nodes of an HA cluster spread across as many system cabinets (Superdome and non-Superdome systems) as possible. This approach maximizes redundancy to further reduce the chance of a failure causing down time.

Windows Server 2003, Datacenter Edition/Enterprise Edition for Itanium 2 systems

Microsoft Cluster Service (MSCS) comes standard with Windows Server 2003. When a customer orders T2372A, Windows Server 2003, Datacenter Edition or T2373A Windows Server 2003, Enterprise Edition for Itanium 2 systems, it includes Microsoft Cluster Service there is no additional SKU or charge for this Windows Server 2003 functionality. MSCS does not come preconfigured from HP's factories, however, so it is recommended that if your customer is interested in a MSCS configuration with Integrity Superdome, HP Services be engaged for a statement of work to configure MSCS on Integrity Superdome with HP storage.

HP Storage is qualified and supported with MSCS clusters. HP storage arrays tested and qualified with MSCS clusters on Superdome are:

- EVA 3000 v3.01
- EVA 5000 v3.01
- XP 48/512
- XP 128/1024.
- XP12000
- MSA1000
- MSA1500

In addition, the following EMC storage arrays are supported with MSCS:

- EMC CLARiiON FC4700
- EMC CLARiiON CX200/CX400/CX600
- EMC CLARiiON CX300/CX500/CX700
- EMC Symmetrix 8000 Family
- EMC DMX 800/1000/2000/3000

HP has qualified and supports the following capabilities with Integrity Superdome and MSCS:

- Active/Active and Active/Passive MSCS clusters
- Partition size: Any size from 2 processors up to 64 processors can be in a cluster
- HP supports anywhere from 2 nodes in an MSCS cluster with Superdome to 8 nodes
- Cluster nodes can be within the same Superdome cabinet or between different Superdome cabinets co located at the same site
- MSCS clusters can be between similar partitions of processor capacity (i.e. 8 processor partition clustered to 8 processor partition, 16 processor partition clustered to 16 processor partition)
- MSCS clusters can also be between dissimilar partitions of processor capacity (i.e. 16 processor partition clustered to 8 processor partition, 32 processor partition clustered to 16 processor partition). Note: You and the customer should work with HP Support to determine the appropriate configuration based on the availability level that is needed by the customer. As an example, if the customer wants a Service Level Agreement based on application availability, then perhaps an exact mirror of the production partition be set up for failover (i.e. similar processor capacity). In any event, please ensure that the proper amount of hardware resources on the target server is available for failover purposes.
- HP Cluster Extention XP is a disaster recovery solution that extends local clusters over metropolitan-wide distance. It now supports MSCS on Windows Integrity with XP48/XP512, XP128/XP1024, XP12000.

For high availability purposes with MSCS, it is recommended (but not required) that customers also use HP SecurePath software (v4.0c SP1) with HP storage for multi pathing and load balancing capabilities in conjunction with the fiber channel HBA, AB466A or AB467A. Additionally, the NCU (NIC Configuration Utility), which is provided from HP on the SmartSetup CD that ships with Windows partitions, can also be used in conjunction with MSCS clusters with the HP supported Windows NIC cards.

Additionally, customers can see the completion of our certification for the Microsoft Windows catalog at the following URL: http://www.microsoft.com/windows/catalog/server/default.aspx?subID=22&xslt=cataloghome&pgn=catalogHome

Microsoft requires hardware vendors to complete this certification - also called "Windows logo-ing."

Configuration

Below is the ordering information for Windows Server 2003 Datacenter Edition (for Itanium 2 based HP Integrity Superdome only).

- To order a system with Windows Server 2003 Datacenter Edition, you must order the T2372A product number with English or Japanese localization (option ABA or ABJ), and the appropriate license to use option code (002 through 064). Windows Server 2003 Datacenter Edition license options should be ordered to accommodate the total number of processors running Windows in the system. Order the fewest option numbers possible for the total license number. For example if there are a total of 24 processors in the system running Datacenter, order options 016 and 008. Datacenter can be partitioned (nPars only) into any number of instances, but is limited to one OS image per nPar.

NOTES:

- Windows Server 2003 Datacenter Edition should be installed with assistance from HP. If factory installation is selected, then a qualified Windows storage device must be ordered and an A9890A or 337972 B21card must be ordered. There must be at least one boot drive for each partition. Two drives are required for RAID 1; one drive is required for RAID 0 or no RAID.
- Cannot order more than one of the same license options.
- Windows Datacenter Edition only supports a maximum of 64 processor coress per partition.
- Quantity of Windows Client Access Licenses with Windows Datacenter Edition, T2372A is 5.

Description	Product/Option Number
Microsoft® Windows® Server 2003 Datacenter Edition for Itanium 2 Systems	T2372A
English Localization	ABA
Japanese Localization	ABJ
Factory Integration	0D1
Include with complete system	B01
2 processor license to use	002
4 processor license to use	004
8 processor license to use	008
16 processor license to use	016
32 processor license to use	032
64 processor license to use	064
HP Standalone Operating System for field install. Windows Server 2003 Datacenter Edition Stand-alone for use when adding licenses to an existing server or replacing another operating system on a Datacenter qualified server. 0D1 (factory integration) is the default operating system installation method and should be used whenever possible. Must be ordered with the appropriate number of licenses (LTU, for example T2372A-016 for 16 Windows Server 2003 Datacenter licenses. There must be a Windows Server 2003 Datacenter Edition processor license for each processor in an Integrity server running Windows. When ordering T2372A-501, the appropriate on-site HP Services installation options will be added to the order.	501

Configuration

Below is the ordering information for Windows Server 2003 Enterprise Edition (for Itanium 2 based HP Integrity Superdome only).
NOTE: All orders for systems using Windows Server 2003 Enterprise Edition (64 bit) must contain T2373A or T2369AA Foundation Pack for Windows LTU but not both. Windows Server 2003, Enterprise Edition is not supported with mx2 modules.

- To order a system with Windows Server 2003 Enterprise Edition: Windows Server 2003 Enterprise Edition supports a maximum of 8 processors per partition; each partition will require a separate instance of the Enterprise Edition OS (T2373A) or a copy of Foundation Pack for Windows LTU (T2369AA). For example, if the Superdome is configured into 4 partitions, the quote must include an order for four (4) T2373As or four (4) T2369AAs or a combination of the two part numbers for a total of four. Configurations can have a maximum of 2 cell boards per partition; a maximum of 1 card cage per cell board (maximum of two card cages per partition) and a maximum of 12 I/O cards per card cage/cell board.
- Configurations requiring greater than 8 processors per partition must select Windows Server 2003 Datacenter Edition (T2372A). Datacenter Edition is ordered on a per processor basis and supports a maximum of 64 processor cores per partition. Moving a partition from Enterprise Edition to Datacenter Edition requires the OS to be reinstalled. If a customer foresees the potential for their environment to grow beyond eight processors per partition it is recommended they purchase Datacenter Edition.
- Quantity of Windows Client Access Licenses with Windows Enterprise Edition, T2373A is 25.

Description	Product/Option Number
Microsoft Windows Server 2003 Enterprise Edition for Itanium 2 Systems	T2373A*
English Localization	ABA
Factory Integration	0D1
Installation and Startup Services are bundled with T2373A. OS support is (optional) for T2373A, however, it is highly recommended that an OS support level is purchased equal to the hardware support level.	
HP Integrity Essential Foundation Pack for Windows	T2369AA
Include with complete system	B01
Installation and Startup Services are required for T2369AA. Please order part number HA114A1-5EK. OS support for T2369AA is a custom quote as the OS is not sold with the system.	
*NOTE: Not orderable/supported in Japan. NOTE: For Japanese versions of Windows Server 2003 Datacenter Edition or Enterprise Edition: The AG064A (the Japanese TFT7600) is supported with Superdome partitions; the AG064A IS NOT SUPPORTED as part of SMS.	

Configuration

Network Adapter Teaming with Windows Server 2003

Windows Server 2003 supports NCU, NIC Configuration Utility. This is the same NCU that is available to Proliant customers. This NCU has been ported to 64 bit Windows Server 2003 and is included with every SmartSetup CD that comes with a Windows partition on Integrity Superdome.

All ProLiant Ethernet network adapters support the following three types of teaming:

- NFT—Network Fault Tolerance
- TLB—Transmit Load Balancing
- SLB—Switch-assisted Load Balancing

For Windows Server 2003 on Superdome, there are four network interface cards that are currently supported (thus, these are the only cards that can be teamed with this NCU):

Windows/Linux PCI 1000Base-T Gigabit Ethernet Adapter (Copper)	A7061A
Windows/Linux PCI 1000Base-SX Gigabit Ethernet Adapter (Fiber)	A7073A
Windows/Linux PCI 2 port 1000Base-T Gigabit Ethernet Adapter (Copper)	A9900A
Windows/Linux PCI 2 port 1000Base-T Gigabit Ethernet Adapter (Fiber)	A9899A

Also, note that teaming between the ports on a single A9900A or A9899A above is supported by the NCU.

OpenVMS Multi-System High Availability

An HP OpenVMS Cluster is a highly integrated organization of AlphaServer and HP Integrity server systems, applications, operating systems, and storage devices. These systems can be connected to each other and storage components in a variety of ways. HP OpenVMS Cluster software is an integral part of the OpenVMS operating system, providing the basis for many of the key capabilities utilized by OpenVMS enterprise solutions. A full "shared everything" cluster design that has been in existence for more than 20 years, OpenVMS Cluster software allows for the maximum in expandability, scalability, and availability for mission-critical applications. Qualified for up to 96 computer nodes and more than 3,000 processors, OpenVMS clusters afford virtually 100 percent uptime and expand the multiprocessing capabilities of the computing environment.

For information on OpenVMS Cluster software, refer to the OpenVMS Cluster Software page http://h71000.www7.hp.com/openvms/products/clusters/index.html and the OpenVMS Cluster Software product description at http://h18000.www1.hp.com/info/SP2978/SP2978PF.PDF. For specific information on supported cluster and storage options for OpenVMS, refer to the OpenVMS Operating System product description at http://h18000.www1.hp.com/info/XAV12X/XAV12XPF.PDF. For information on specific Cluster Configuration guidelines, refer to the OpenVMS documentation at http://h71000.www7.hp.com/doc/, specifically the latest version of the Guidelines to OpenVMS Cluster Configuration manual.

Management Features

Integrity Superdome now supports the Console and Support Management Station in one device.

Configuration

Service Processor (MP)

The service processor (MP) utility hardware is an independent support system for nPartition servers. It provides a way for you to connect to a server complex and perform administration or monitoring tasks for the server hardware and its nPartitions. The main features of the service processor include the Command menu, nPartition consoles, console logs, chassis code viewers, and partition Virtual Front Panels (live displays of nPartition and cell states).

Access to the MP is restricted by user accounts. Each user account is password protected and provides a specific level of access to the Superdome complex and service processor commands. Multiple users can independently interact with the service processor because each service processor login session is private. Up to 16 users can simultaneously log in to the service processor through its network (customer LAN) interface and they can independently manage nPartitions or view the server complex hardware states. Two additional service processor login sessions can be supported by the local and remote serial ports. These allow for serial port terminal access(through the local RS-232 port) and external modem access (through the
remote RS-232 port).

In general, the service processor (MP) on Superdome servers is similar to the service processor on other HP servers, while providing enhanced features necessary for managing a multiple-nPartition server. For example, the service processor manages the complex profile, which defines nPartition configurations as well as complex-wide settings for the server. The service processor also controls power, reset, and TOC capabilities, displays and records system events (chassis codes), and can display detailed information about the various internal subsystems.

Functional capabilities

The primary features available through the service processor are:

- The Service Processor Command Menu: provides commands for system service, status, access configuration, and manufacturing tasks.
- Partition Consoles: Each nPartition in a server complex has its own console. Each nPartition's console provides access to Boot Console Handler (BCH) interface and the HP-UX console for the nPartition.
- Console Logs: Each nPartition has its own console log, which has a history of the nPartition console's output, including boot output, BCH activity, and any HP-UX console login activity.
- Chassis Logs Viewers (Live and Recorded Chassis Codes): Three types of chassis code log views are available: activity logs, error logs, and live chassis code logs.
- Virtual Front Panels: Each nPartition's Virtual Front Panel (VFP) displays real-time status of the nPartition boot status and activity, and details about all cells assigned to the nPartition. The VFP display automatically updates as cell and nPartition status changes.

Configuration

Support Management Station (SMS)

The Support Management Station (SMS) runs the Superdome scan tools that enhance the diagnosis and testability of Superdome. The SMS and associated tools also provide for faster and easier upgrades and hardware replacement.

The purpose of the SMS is to provide Customer Engineers with an industry leading set of support tools, and thereby enable faster troubleshooting and more precise problem root cause analysis. It also enables remote support by factory experts who consult with and back up the HP Customer Engineer. The SMS complements the proactive role of HP's Instant Support Enterprise Edition (ISEE) that is offered to Mission Critical customers by focusing on reactive diagnosis for both mission critical and non mission critical Superdome customers.

The user of the SMS is the HP Customer Engineer and HP Factory Support Engineer. The Superdome customer benefits from their use of the SMS by receiving faster return to normal operation of their Superdome server and improved accuracy of fault diagnosis, resulting in fewer callbacks. HP can offer better service through reduced installation time.

Functional Capabilities: The SMS basic functional capabilities are:

- Remote access via customer LAN
- Modem access (PA-8800/PA-8900 SMS only)
- Ability to be disconnected from the Superdome platform(s) and not disrupt their operation.
- Ability to connect a new Superdome platform to the SMS and be recognized by scan software.
- Support for up to sixteen Superdome systems
- Ability to support multiple, heterogeneous Superdome platforms (scan software capability).
- System scan and diagnostics
- Utility firmware updates
- Enhanced IPMI logging capabilities (Windows-based ProLlant SMS only)

Console Access

The optimal configuration of console device(s) depends on a number of factors, including the customer's data center layout, console security needs, customer engineer access needs, and the degree with which an operator must interact with server or peripheral hardware and a partition (i.e. changing disks, tapes). This section provides a few guidelines. However the configuration that makes best sense should be designed as part of site preparation, after consulting with the customer's system administration staff and the field engineering staff.

Customer data centers exhibit a wide range of configurations in terms of the preferred physical location of the console device. (The term "console device" refers to the physical screen/keyboard/mouse that administrators and field engineers use to access and control the server.) The Superdome server enables many different configurations by its flexible configuration of access to the MP, and by its support for multiple geographically distributed console devices.

Three common data center styles are:

- The secure site where both the system and its console are physically secured in a small area.
- The "glass room" configuration where all the systems' consoles are clustered in a location physically near the machine room.
- The geographically dispersed site, where operators administer systems from consoles in remote offices.

These can each drive different solutions to the console access requirement.

The considerations listed below apply to the design of provision of console access to the server. These must be considered during site preparation.

- The Superdome server can be operated from a VT100 or an hpterm compatible terminal emulator. However some programs (including some of those used by field engineers) have a more friendly user interface when operated from an hpterm.
- LAN console device users connect to the MP (and thence to the console) using terminal emulators that establish telnet connections to the MP. The console device(s) can be anywhere on the network connected to either port of the MP.
- Telnet data is sent between the client console device and the MP "in the clear", i.e. unencrypted. This may be a concern for some customers, and may dictate special LAN configurations.
- If an HP-UX workstation is used as a console device, an hpterm window running telnet is the recommended way to connect to the MP. If a PC is used as a console device, Reflection1 configured for hpterm emulation and telnet connection is the recommended way to connect to the MP.
- The MP currently supports a maximum of 16 telnet-connected users at any one time.
- It is desirable, and sometimes essential for rapid time to repair to provide a reliable way to get console access that is physically close to the server, so that someone working on the server hardware can get immediate access to the results of their actions. There are a few options to achieve this:
  - Place a console device close to the server.
  - Ask the field engineer to carry in a laptop, or to walk to the operations center.
  - Use a system that is already in close proximity of the server such as the Instant Support Enterprise Edition (ISEE) or the System Management Station as a console device close to the system.
  - The system administrator is likely to want to run X-applications or a browser using the same client that they access the MP and partition consoles with. This is because the partition configuration tool, *parmgr*, has a graphical interface. The system administrator's console device(s) should have X-window or browser capability, and should be connected to the system LAN of one or more partitions.

Functional capabilities

- Local console physical connection (RS-232)
- Display of system status on the console (Front panel display messages)
- Console mirroring between LAN and RS-232 ports
- System hard and soft (TOC or INIT) reset capability from the console.
- Password secured access to the console functionality
- Support of generic terminals (i.e. VT100 compatible).
- Power supply control and monitoring from the console. It will be possible to get power supply status and to switch power on/off from the console.
- Console over the LAN. This means that a PC or HP workstation can become the system console if properly connected on the customer LAN. This feature becomes especially important because of the remote power management capability. The LAN will be implemented on a separate port, distinct from the system LAN, and provide TCP/IP and Telnet access.
- There is one MP per Superdome cabinet, thus there are two (2) for Superdome 64-processor. But one, and only one, can be active at a time. There is no redundancy or failover feature.

Configuration

Windows Server 2003

For Windows Server 2003 customers desiring uninterrupted visibility to the Superdome Windows partition, it is recommended that customers purchase an IP console solution separately to view the partition while the OS is rebooting (in addition to normal Windows desktop if desired). Windows Terminal Services (standard in Windows Server 2003) should be the recommended method to provide remote access, but is lacking in displaying VGA output during reboot.

For customers who mandate VGA access during reboot, the IP console switch (262586-B21), used in conjunction with a VGA/USB card in the partition (A6869A/A6869B) is the solution. These IP console solutions are available "off the shelf" with resellers or the ProLiant supply chain.

The features of this switch are as follows:

- Provides keyboard, video and mouse (KVM) connections to 16 direct attached Windows partitions (or servers) - expandable to 128.
- Allows access to partitions (or servers) from a remote centralized console.
- 1 for local KVM
- 3 concurrent remote users (secure SSL data transfer across network)
- Single screen switch management with the IP Console Viewer Software:
  - Authentication
  - Administration
  - Client Software

If the full graphical console access is needed, the following must be ordered with the Integrity Superdome purchase (it will not be integrated in the factory, but will ship with the system):

Component	Product Number
3×1×16 IP console switch (100 240V)-1 switch per 16 OS instances (n<=16), each connected to VGA card	262586-B21
8 to 1 console expander-Order expander if there are more than 16 OS instances	262589-B21
USB interface adapters-Order one per OS instance	336047-B21
CAT5 cable-Order one per OS instance	

NOTE: For Japanese versions of Windows Server 2003 Datacenter Edition or Enterprise Edition: The (the Japanese TFT7600) is supported with Superdome partitions; the AG064A IS NOT SUPPORTED as part of SMS.

For additional information, please visit:
http://h18004.www1.hp.com/products/servers/proliantstorage/rack-options/kvm/index-console.html

Configuration

Support

The following matrix describes the supported SMS and recommended console devices for all Superdomes.

SMS and Console Support Matrix

	SMS	Console
PA 8700 (pre March 1, 2004)	Legacy UNIX SMS [1]	PC/workstation
PA 8700 (post March 1, 2004)	UNIX rx2600 bundle	TFT5600 + Ethernet switch
PA 8700 upgraded to Integrity or PA-8800/PA-8900	Legacy UNIX SMS with software upgrade [2]	PC/workstation
	UNIX rx2600 bundle [3]	TFT5600 + Ethernet switch
	Windows SMS/Console (ProLiant ML350)	
Integrity or PA-8800/PA-8900	Windows SMS/Console (ProLiant ML350)	
	UNIX SMS/Console (rx2600)	

[1] A legacy UNIX SMS could be an A400, A500, rp2430 or rp2470 bundle, depending on when it was ordered

[2] In order for a legacy SMS to be upgraded to support Integrity or PA8800/PA-8900, it must be running HP-UX 11.0 or later, as sx1000 scan tools are not supported on HP-UX 10.20.

[3] rx2600 SMS bundles ordered and installed prior to October 2004 will require a software upgrade in order to support an sx1000-based Superdome. As of October 2004, all rx2600 SMS bundles support PA8800/PA-8900 and Integrity Superdomes without this upgrade.

Hardware Requirements

There are two alternatives for Integrity Superdome: a Windows ProLiant SMS/Console or a UNIX rx2600 SMS/Console.

The UNIX rx2600 SMS bundle is comprised of:

- HP rx2600 1.0G 1.5MB processor server Solution
- Factory rack kit for rx2600
- 1GB DDR memory
- 36GB 15K HotPlug Ultra320 HDD
- HP-UX 11iv2 Foundation OE
- 1 x HP Tape Array 5300 with DVD-ROM and DAT 40
- HP ProCurve Switch 2124
- CAT 5e Cables

By default, the rx2600 SMS does not come with a display monitor or keyboard unless explicitly ordered to enable console access (the TFT5600 rackmounted display/mouse/keyboard is the recommended solution). See the ordering guide for details on the additional components that are required in order to use the rx2600 SMS as a console.

The Windows ProLiant SMS:

- Allows local access to SMS by CE.
- Provides integrated console access, providing hpterm emulation over telnet and web browser, connecting over LAN or serial to a Superdome system
- Provides remote access over a LAN or dialup connection:
- ftp server with capability to ftp the firmware files and logs
- dialup modem access support (i.e., PC Anywhere or VNC)
- Provides seamless integration with data center level management.
- Provides partition logon capability, providing hpterm emulation over telnet, X windows, and Windows Terminal Services capabilities.
- Provides following diagnostics tools:

- Runs HP's proven highly effective JTAG scan diagnostic tools, which offer rapid fault resolution to the failing wire.
- Console log storage and viewing
- Event log storage and viewing
- Partition and memory adviser flash applications
- Supports updating platform and system firmware.
- Always on event and console logging for Superdome systems, which captures and stores very long event and console histories, and allows HP specialists to analyze the first occurrence of a problem.
- Allows more than one LAN connected response center engineer to look at SMS logs simultaneously.
- Can be disconnected from the Superdome systems and not disrupt their operation.
- Provides ability to connect a new Superdome system to the SMS and be recognized by scan software.
- Scans one Superdome system while other Superdome systems are connected (and not disrupt the operational systems).
- Supports multiple, heterogeneous Superdome platforms.

The Windows ProLiant SMS/Console is comprised of a ProLiant ML350 G3/G4 and a TFT5600 retractable display monitor/keyboard/mouse to enable console access. This solution also requires (and includes) a switch. This is because scan diagnostics will not work properly if more than one IP address exists on the ProLiant SMS/Console. An important difference between the UNIX rx2600 SMS and the Windows ProLiant SMS/Console is that the ProLiant SMS, by default, provides console functionality whereas the UNIX rx2600 SMS does not. The TFT5600 retractable display/keyboard which is an optional add-on for the UNIX-based rx2600 SMS, is included by default with the Windows ProLiant SMS.

Additionally, the Windows ProLiant SMS/Console includes an internal modem that is intended for connection to a phone line. This is for cases in which the Customer does not want the SMS to be on a public network, and HP Field Services needs to access the SMS (they would then access the SMS via the phone line and PCAnywhere.)

A customer may not substitute any PC running Windows Server 2000 SP4 for the ProLiant SMS/Console due to the specialized software applications that have been qualified on the hardware and OS. Utilizing any other device as the SMS will void the warranty on the Superdome system and degrade the ability to service the system.

The Windows ProLiant SMS bundle is comprised of:

- 1 x HP ProLiant server ML350 G3/G4
- 1 x Intel Xeon ™ DP 3.06 GHz processor
- 2 x 256MB
- 36GB 10K U320 HDD
- 1 x internal DVD
- 1 x internal V.90 56K modem with phone cord
- Windows 2000 Server SP4
- 2 x 25' CAT5e cables
- 1 x 4' CAT5e cable
- Ethernet switch and jumper cord
- Retractable display/keyboard/mouse
- Third-party applications

NOTE: If full graphical access to the SMS is needed, the PS/2 Interface Adapter (262588 B21) will allow the SMS to share the IP Console Switch with other OS instances.

Configuration

Software Requirements

All SMS software is preloaded in the factory and delivered to the customer as a complete solution.

The UNIX rx2600 SMS supports only HP UX 11i at this time. Current versions of the SMS software have not been qualified for 64 bit Windows.

The Windows ProLiant SMS/Console will run Windows 2000 SP4 as the default operating system. The ProLiant SMS/Console will follow the Windows OS roadmap and support later versions of this operating system as needed. The version of the scan tools used on the sx1000 Superdomes also does not require scan traffic to be isolated from console traffic.

SMS/Console Connectivity

The UNIX rx2600 SMS requires the addition of a retractable display/keyboard in order to provide console access to Integrity Superdome.

Console capabilities are integrated with the Windows ProLiant SMS solution. The ProLiant SMS includes the necessary display, keyboard and mouse, and only one LAN port on the MP is required to be connected. The Windows ProLiant SMS can support and connect to either of the MP LAN interfaces (known as the Private and Customer LAN ports). Both LAN ports on the MP have identical functionality so there is no preference in using one over the other. Only one IP port on the Superdome MP is required to be connected to the ProLiant SMS. Since scan and console traffic can co-exist on the same network, only one IP address exists (and is supported) on the ProLiant SMS.

The Integrity Superdome scan tools use TCP/IP, not UDP, for scan diagnostics. Therefore it is not necessary to isolate scan and console traffic on Integrity Superdomes.

The Core I/O cards from each nPar can optionally be connected to the Ethernet switch in order to facilitate graphical console functionality (i.e., parmgr). However, security concerns may dictate that a partition NIC not be connected to the management LAN. Alternatives are to access from a management station to a partition LAN through a secure router, or to use text mode access to commands via the console.

Additional Integrity Superdomes can be added to this configuration. Each new Integrity Superdome will require only one CAT5e LAN cable for connection of the Customer/Private LAN port on the Superdome MP to the existing switch.

Customers using the UNIX rx2600 SMS solution for Integrity Superdome can also use the rx2600 as a console device. In order to use the UNIX rx2600 SMS as a console for Integrity Superdome, the only component that must be ordered is the TFT5600 retractable keyboard/display/mouse. See the ordering guide for more details.

Additional Integrity Superdomes require only a CAT5e LAN cable for connection of the Private/Customer LAN port to the existing switch. PA-8700 Superdomes can also be added to this configuration, but require separate scan & console networks. Details on mixed Superdome environments follow later in this section.

Windows Terminal Services (standard in Windows Server 2003) should be the recommended method to provide remote access for Windows partitions, but is lacking in displaying VGA output during reboot. For customers who mandate VGA access during reboot, configure the IP console switch (262586 B21) used in conjunction with a VGA/USB card in the partition (A6869A).

Configuration

Mixed Superdome Environments

Some important rules regarding SMS/Console support in mixed Superdome environments:

1. PA-8700 Superdomes require scan and console traffic to be isolated on separate networks
2. Integrity Superdomes do not require scan and console traffic to be isolated on separate networks
3. The Windows ProLiant SMS cannot be used to manage PA-8700 Superdomes
4. The UNIX rx2600 SMS can be used to manage all current models of Superdome
5. A legacy UNIX SMS can manage Integrity Superdomes as long as it undergoes a software upgrade and is running HP-UX 11.0

Upgrading a PA-8700 Superdome to Integrity

Upgrading a PA-8700 Superdome to Integrity
Once a PA-8700 Superdome has been upgraded to Integrity, a software upgrade must be performed on the SMS in order to have the correct scan tools for the sx1000 chipset. Alternatively, a new Windows ProLiant SMS or a UNIX rx2600 SMS can be ordered.

After performing the software upgrade, scan and console traffic from the SMS to the newly upgraded Superdome can co-exist on the same network. Note that if there are other PA-8700 Superdomes still connected to the SMS, those Superdomes will still require the two separate networks (see Adding a PA-8800/PA-8900 Superdome to a PA-8700 Environment, below).

Adding an Integrity Superdome to a PA-8700 Environment

When adding an Integrity Superdome to a PA-8700 environment, the Customer can choose to purchase a new SMS/Console (either the Windows or UNIX version) to manage the new Integrity Superdome, and continue to use their existing legacy SMS to manage the PA-8700 Superdomes.

A simpler solution is to upgrade the software on the legacy UNIX SMS so that it may manage the new Integrity Superdome. The upgraded UNIX SMS is able to determine which Superdomes are PA-8700 and which are Integrity, and will separate scan traffic from console traffic on the PA-8700 Superdomes.

To support this, the separate scan and console networks should be maintained for all PA-8700 Superdomes. The new Integrity Superdome requires only one connection to the SMS from its Private/Customer LAN port on the MP.

System Management Features

HP-UX (sx1000 & sx2000)

- HP-UX Servicecontrol Manager is the central point of administration for management applications that address the configuration, fault, and workload management requirements of an adaptive infrastructure.
- Servicecontrol Manager maintains both effective and efficient management of computing resources. It integrates with many other HP-UX-specific system management tools, including the following, which are available on Itanium 2 based servers:
- Ignite-UX-Ignite-UX addresses the need for HP-UX system administrators to perform fast deployment for one or many servers. It provides the means for creating and reusing standard system configurations, enables replication of systems, permits post-installation customizations, and is capable of both interactive and unattended operating modes.
- Software Distributor (SD) is the HP-UX administration tool set used to deliver and maintain HP-UX operating systems and layered software applications. Delivered as part of HP-UX, SD can help you manage your HP-UX operating system, patches, and application software on HP Itanium2-based servers.
- System Administration Manager (SAM) is used to manage accounts for users and groups, perform auditing and security, and handle disk and file system management and peripheral device management. Servicecontrol Manager enables these tasks to be distributed to multiple systems

and delegated using role based security.

- HP-UX Kernel Configuration-for self-optimizing kernel changes. The new HP-UX Kernel Configuration tool allows users to tune both dynamic and static kernel parameters quickly and easily from a Web based GUI to optimize system performance. This tool also sets kernel parameter alarms that notify you when system usage levels exceed thresholds.
- Partition Manager creates and manages nPartitions-hard partitions for high-end servers. Once the partitions are created, the systems running on those partitions can be managed consistently with all the other tools integrated into Servicecontrol Manager. Key features include:
  - Easy-to-use, familiar graphical user interface.
  - Runs locally on a partition, or remotely. The Partition Manager application can be run remotely on any system running HP-UX 11i Version 2 and eventually select Windows releases and remotely manage a complex either by 1) communicating with a booted OS on an nPartition in the target complex via WBEM, or 2) communicating with the service processor in the target complex via IPMI over LAN. The latter is especially significant because a complex can be managed with NONE of the nPartitions booted.
  - Full support for creating, modifying, and deleting hardware partitions.
  - Automatic detection of configuration and hardware problems.
  - Ability to view and print hardware inventory and status.
  - Big picture views that allow system administrators to graphically view the resources in a server and the partitions that the resources are assigned to.
  - Complete interface for the addition and replacement of PCI devices.
  - Comprehensive online help system.
- Security Patch Check determines how current a systems security patches are, recommends patches for continuing security vulnerabilities and warns administrators about recalled patches still present on the system.
- System Inventory Manager is for change and asset management. It allows you to easily collect, store and manage inventory and configuration information for HP-UX based servers. It provides an easy-to-use, Web-based interface, superior performance, and comprehensive reporting capabilities
- Event Monitoring Service (EMS) keeps the administrator of multiple systems aware of system operation throughout the cluster, and notifies the administrator of potential hardware or software problems before they occur. HP Servicecontrol Manager can launch the EMS interface and configure EMS monitors for any node or node group that belongs to the cluster, resulting in increased reliability and reduced downtime.
- Process Resource Manager (PRM) controls the resources that processes use during peak system load. PRM can manage the allocation of processor, memory resources, and disk bandwidth. It allows administrators to run multiple mission critical applications on a single system, improve response time for critical users and applications, allocate resources on shared servers based on departmental budget contributions, provide applications with total resource isolation, and dynamically change configuration at any time-even under load. (fee based)
- HP-UX Workload Manager (WLM) A key differentiator in the HP-UX family of management tools, Workload Manager provides automatic processor resource allocation and application performance management based on prioritized service-level objectives (SLOs). In addition, WLM allows administrators to set real memory and disk bandwidth entitlements (guaranteed minimums) to fixed levels in the configuration. The use of workload groups and SLOs improves response time for critical users, allows system consolidation, and helps manage user expectations for performance. (Fee-based)
- HP's Management Processor enables remote server management over the Web regardless of the system state. In the unlikely event that none of the nPartitions are booted, the Management Processor can be accessed to power cycle the server, view event logs and status logs, enable console redirection, and more. The Management Processor is embedded into the server and does not take a PCI slot. And, because secure access to the Management Processor is available through SSL encryption, customers can be confident that its powerful capabilities will be available only to authorized administrators. New features that will be available include:

  - Support for Web Console that provides secure text mode access to the management processor
  - Reporting of error events from system firmware.
  - Ability to trigger the task of PCI OL* from the management processor.
  - Ability to scan a cell board while the system is running. (only available for partitionable systems)
  - Implementation of management processor commands for security across partitions so that partitions do not modify system configuration (only available for partitionable systems).
- OpenView Operations Agent-collects and correlates OS and application events (fee based)
- OpenView Performance Agent-determines OS and application performance trends (fee based)
- OpenView GlancePlus-shows real time OS and application availability and performance data to diagnose problems (fee based)
- OpenView Data Protector (Omniback II)-backs up and recovers data (fee based)

In addition, the Network Node Manager (NNM) management station will run on HP-UX Itanium 2 based servers. NNM automatically discovers, draws (maps), and monitors networks and the systems connected to them.

All other OpenView management tools, such as OpenView Operations, Service Desk, and Service Reporter, will be able to collect and process information from the agents running on Itanium 2-based servers running HP-UX.

Windows Server 2003, Datacenter Edition/Enterprise Edition (sx1000 and sx2000)

- The HP Essentials Foundation Pack for Windows is a complete toolset to install, configure, and manage Itanium2 servers running Windows. Included in the Pack is the Smart Setup DVD which contains all the latest tested and compatible HP Windows drivers, HP firmware, HP Windows utilities, and HP management agents that assist in the server deployment process by preparing the server for installation of standard Windows operating system and in the on going management of the server. Please note that this is available for HP service personnel but not provided to end customers.
- Partition Manager and Partition Commands
  Two tools are available for managing nPartition (hard partition) configurations on Superdome. The first tool is the Partition Manager (parmgr) is a web-based application that allows system administrators to use a convenient graphical user interface (GUI) to configure and manage nPartitions on HP server systems. Partition Manager can also detect several types of configuration problems. Partition Manager can be accessed on the SMS through the System Management Homepage (see below).
  The second tool is Partition Commands (parCLI) is a set of command-line utilities for viewing and modifying the nPartition configuration of HP server systems. All functionality provided by the Partition Commands are
  also available via the Partition Manager GUI, however the commands may be used for scripting common tasks. The Partition Commands may be run on the SMS from a Command Prompt window or batch script.
  Key features of both tools include full support for creating, modifying, and deleting hardware partitions. Refer to the HP UX section above for additional features of Partition Manager and the Partition Commands.
- Insight Manager 7 maximizes system uptime and provides powerful monitoring and control. Insight Manager 7 delivers pre-failure alerting for servers ensuring potential server failures are detected before they result in unplanned system downtime. Insight Manager 7 also provides inventory reporting capabilities that dramatically reduce the time and effort required to track server assets and helps systems administrators make educated decisions about which system may required hardware upgrades or replacement. And Insight Manager 7 is an effective tool for managing your HP desktops and notebooks as well as non HP devices instrumented to SNMP or DMI.

- System Management Homepage displays critical management information through a simple, task oriented user interface. All system faults and major subsystem status are now reported within the initial System Management Homepage view. In addition, the new tab-based interface and menu structure provide one click access to server log. The System Management Homepage is accessible either directly through a browser (with the partition's IP address) or through a management application such as Insight Manager 7 or an enterprise management application.
- HP's Management Processor enables remote server management over the Web regardless of the system state. In the unlikely event that the operating system is not running, the Management Processor can be accessed to power cycle the server, view event logs and status logs, enable console redirection, and more. The Management Processor is embedded into the server and does not take a PCI slot. And, because secure access to the Management Processor is available through SSL encryption, customers can be confident that its powerful capabilities will be available only to authorized administrators. New features on the management processor include:
  - Support for Web Console that provides secure text mode access to the management processor
  - Reporting of error events from system firmware.
  - Ability to trigger the task of PCI OL* from the management processor.
  - Ability to scan a cell board while the system is running.
  - Implementation of management processor commands for security across partitions so that partitions do not modify system configuration.
- OpenView Management Tools, such as OpenView Operations and Network Node Manager, will be able to collect and process information from the SNMP agents and WMI running on Windows Itanium 2 based servers. In the future, OpenView agents will be able to directly collect and correlate event, storage, and performance data from Windows Itanium 2-based servers, thus enhancing the information OpenView management tools will process and present.

Red Hat RHEL AS 3 & 4 and SUSE SLES 9 (sx1000 & sx2000)

- Insight Manager 7 maximizes system uptime and provides powerful monitoring and control. Insight Manager 7 also provides inventory reporting capabilities that dramatically reduce the time and effort required to track server assets and helps systems administrators make educated decisions about which system may required hardware upgrades or replacement. And Insight Manager 7 is an effective tool for managing your HP desktops and notebooks as well as non HP devices instrumented to SNMP or DMI.
- The HP Enablement Kit for Linux facilitates setup and configuration of the operating system. This kit includes System Imager, an open source operating system deployment tool. System Imager is a golden image based tool and can be used for initial deployment as well as updates.
- Partition Manager creates and manages nPartitions-hard partitions for high-end servers. Once the partitions are created, the systems running on those partitions can be managed consistently with all the other tools integrated into Servicecontrol Manager.
  NOTE: At first release, Partition Manager will require an HP-UX 11i Version 2 partition or separate device (i.e. Itanium2 based workstation or server running HP-UX 11i Version 2) in order to configure Red Hat or SUSE partitions. Refer to HP UX section above for key features of Partition Manager.
- HP's Management Processor enables remote server management over the Web regardless of the system state. In the unlikely event that the operating system is not running, the Management Processor can be accessed to power cycle the server, view event logs and status logs, enable console redirection, and more. The Management Processor is embedded into the server and does not take a PCI slot. And, because secure access to the Management Processor is available through SSL encryption, customers can be confident that its powerful capabilities will be available only to authorized administrators.
  - Support for Web Console that provides secure text mode access to the management processor
  - Reporting of error events from system firmware.

- Ability to trigger the task of PCI OL* from the management processor.
  NOTE: Online addition/replacement (OLAR) is not supported when running Red Hat or SUSE in the partition.
- Ability to scan a cell board while the system is running. (only available for partitionable systems)
- Implementation of management processor commands for security across partitions so that partitions do not modify system configuration. (only available for partitionable systems)

OpenVMS
Software Deployment

- Factory Installed Software

Configuration

- Partition Manager creates and manages nPartitions-hard partitions for high?end servers. Once the partitions are created, the systems running on those partitions can be managed consistently with all the other tools integrated into HP Systems Insight Manager. See "Partitioning" for more information.
- OpenVMS Management Station to manage user accounts, printers, and disks
- Availability Manager for real-time performance monitoring

Workload Management

- Global Workload Manager (gWLM) - Global Workload Manager provides automatic CPU resource allocation and application performance management based on prioritized service level objectives (SLOs).
- Class Scheduler for resource management

System Management for OpenVMS

- HP Systems Insight Manager (see above) in conjunction with (Web) Management Agents
- Central Management Server - CMS - Management agent for gWLM

OpenView for OpenVMS

- OpenView Operations Agent-collects and correlates OS and application events (fee-based)
- OpenView Performance Agent -- determines OS and application performance trends (fee-based)

Configuration

General Site Preparation Rules

AC Power Requirements

The modular, N+1 power shelf assembly is called the Front End Power Subsystem (FEPS). The redundancy of the FEPS is achieved with 6 internal Bulk Power Supplies (BPS), any five of which can support the load and performance requirements.

Input Options

Reference the Site Preparation Guide for detailed power configuration options.

Input Power Options

PDCA Product Number	Source Type	Source Voltage (nominal)	PDCA Required	Input Current Per Phase 200-240 VAC	Power Required
A5800A Option 006	3-phase	Voltage range 200-240 VAC, phase-to-phase, 50/60 Hz	4-wire	44 A Maximum per phase	2.5 meter UL power cord and OL approved plug provided. The customer must provide the mating in line connector or purchase quantity one A6440A opt 401 to receive a mating in line connector. An electrician must hardwire the in-line connector to 60 A/63 A site power.[a,b,c]
A5800A Option 007	3-phase	Voltage range 200-240 VAC, phase-to-neutral, 50/60 Hz	5-wire	24 A Maximum per phase	2.5 meter <HAR> power cord and VDE approved plug provided. The customer must provide the mating in line connector or purchase quantity 1 A6440A opt 501 to receive a mating in-line connector. An electrician must hardwire the in-line connector to 30 A/32 A site power.[a,b,d]

a. A dedicated branch is required for each PDCA installed.
b. In the U.S.A, site power is 60 Amps; in Europe site power is 63 Amps.
c. Refer to the Option 006 and 007 Specifics Table for detailed specifics related to this option.
d. In the U.S.A. site power is 30 Amps; in Europe site power is 32 Amps.

Option 006 and 007 Specifics*

			Customer Provided Part	
PDCA Product Number	Attached Power Cord	Attached Plug	In Line Connector	Panel Mount Receptacle
A5800A Option 006	OLFLEX 190 (PN 600804), four-conductor, 6-AWG (16 mm^2), 600-Volt, 60-Amp, 90-degree C, UL, and CSA approved, conforms to CE directives GN/YW ground wire.	Mennekes ME 460P9 3-phase, 4-wire, 60-Amp, 250-Volt, UL-approved. Color blue, IEC 309-1, IEC 309-1, grounded at 3:00 o'clock.	Mennekes ME 460C9 3-phase, 4-wire, 60-amp, 250-Volt, UL-approved. Color blue, IEC 309-1, IEC 309-1, grounded at 9:00 o'clock.[a]	Mennekes ME 460R9 3 phase, 4-wire, 60-amp, 250-Volt, UL-approved. Color blue, IEC 309-1, IEC 309-1, grounded at 9:00 o'clock.[b]
A5800A Option 007	Five conductors, 10-AWG (6 mm^2), 450/475-Volt, 32-Amp, <HAR> European wire cordage, GN/YW ground wire.	Mennekes ME 532P6-14 3-phase, 5-wire, 32-Amp, 450/475-volt, VDE-certified, color red, IEC 309-1, IEC 309-2, grounded at 6:00 o'clock.	Mennekes ME 532C6 16 3-phase, 5-wire, 32-Amp, 450/475-Volt, VDE-certified, color red, IEC 309-1, IEC 309-2, grounded at 6:00 o'clock.[c]	Mennekes ME532R6 1276 3-phase, 5-wire, 32-Amp, 450/475-Volt, VDE- certified, color red, IEC 309-1, IEC 309-2, grounded at 6:00 o'clock.[b]

a. In-line connector is available from HP by purchasing A6440A, Option 401.
b. Panel mount receptacles must be purchased by the customer from a local Mennekes supplier.
c. In-line connector is available from HP by purchasing A6440A, Option 501.

NOTE: A qualified electrician must wire the PDCA in line connector to site power using copper wire and in compliance with all local codes.

Input Requirements

Reference the Site Preparation Guide for detailed power configuration requirements.

Requirements	Value	Conditions/Comments
Nominal Input Voltage (VAC rms)	200/208/220/230/240	
Input Voltage Range (VAC rms)	200-240	Auto-selecting. Measure at input terminals
Frequency Range (Hz)	50/60	
Number of Phases	3	3-phase 5-wire with power cord; 3-phase 4-wire with power cord
Maximum Input Current (A rms), 3-Phase 5-wire	20	3-phase source with a source voltage of 220 VAC measured phase to neutral
Maximum Input Current (A rms), 3-Phase 4-wire	40	3-phase source with a source voltage of either 208 VAC or 230 VAC measured phase to phase
Maximum Inrush Current (A peak)	90	
Circuit Breaker Rating (A), 3-Phase 5-wire	25 A	Per phase
Circuit Breaker Rating (A), 3-Phase 4-wire	45 A	Per phase
Power Factor Correction	0.95 minimum	

Ground Leakage Current (mA)	>3.5 mA, with 6 BPSs installed	Warning label applied to the PDCA at the AC Mains input

Cooling Requirements

- The cooling system in Superdome was designed to maintain reliable operation of the system in the specified environment. In addition, the system is designed to provide redundant cooling (i.e. N+1 fans and blowers) that allows all of the cooling components to be "hot swapped."
- Superdome was designed to operate in all data center environments with any traditional room cooling scheme (i.e. raised floor environments) but in some cases where data centers have previously installed high power density systems, alternative cooling solutions may need to be explored by the customer. HP has teamed with Liebert to develop an innovative data room cooling solution called DataCool. DataCool is a patented overhead climate system utilizing fluid based cooling coils and localized blowers capable of cooling heat loads of several hundred watts per square foot. Some of DataCool's highlights are listed below:
- Liebert has filed for several patents on DataCool
- DataCool, based on Liebert's TeleCool, is an innovative approach to data room cooling
- Liquid cooling heat exchangers provide distributed cooling at the point of use
- Delivers even cooling throughout the data center preventing hot spots
- Capable of high heat removal rates (500 W per square foot)
- Floor space occupied by traditional cooling systems becomes available for revenue generating equipment.
- Enables cooling upgrades when installed in data rooms equipped with raised floor cooling

DataCool is a custom-engineered overhead solution for both new data center construction and for data room upgrades for high heat loads. It is based on Liebert's TeleCool product, which has been installed in 600 telecommunications equipment rooms throughout the world. The system utilizes heat exchanger pump units to distribute fluid in a closed system through patented cooling coils throughout the data center. The overhead cooling coils are highly efficient heat exchangers with blowers that direct the cooling where it is needed. The blowers are adjustable to allow flexibility for changing equipment placement or room configurations. Equipment is protected from possible leaks in the cooling coils by the patented monitoring system and purge function that detects any leak and safely purges all fluid from the affected coils. DataCool has interleaved cooling coils to enable the system to withstand a single point of failure and maintain cooling capability.

Features and Benefits

- Fully distributed cooling with localized distribution
- Even cooling over long distances
- High heat load cooling capacity (up to 500 W per square foot)
- Meets demand for narrow operating temperature for computing systems
- Allows computer equipment upgrade for existing floor cooled data rooms
- Floor space savings from removal of centralized air distribution
- Withstand single point of failures

For More Information

http://www.liebert.com/assets/products/english/products/env/datacool/60hz/bro_8pg/acrobat/sl_16700.pdf

HP has entered into an agreement with Liebert to reference sell the DataCool solution

- The HP/Liebert business relationship is managed by the HP Complementary Products Division.
- DataCool will be reference by HP. Liebert will perform installation, service and support.
- HP will compensate the HP Sales Representative and District Manager for each DataCool that Liebert sells to a customer referred by HP.
- An HP/Liebert DataCool website will be setup to get more information on the product and to manage the reference sales process. Please go to http://hpcp.grenoble.hp.com/ for more information.

Environmental

- 68 to 86 degrees F (20 to 30 degrees C) inlet ambient temperature
- 0 to 10,000 feet (0 to 3048 meters)
- 2600 CFM with N+1 blowers. 2250 CFM with N.
- 65 dBA noise level

Uninterruptible Power Supplies (UPS)

HP will be reselling-high end (10 kW and above) three-phase UPS systems from our partners. We will test and qualify a three-phase UPS for Superdome.

- All third-party UPS resold by HP are tested and qualified by HP to ensure interoperability with our systems
- We plan to include ups_mond ups communications capability in the third party UPS(s), thus ensuring consistent communications strategy with our PowerTrust UPS(s)
- We will also establish a support strategy with our third-party UPS partners to ensure the appropriate level of support our customer have come to expect from HP.
- For more information on the product and to manage the reference sales process please go to http://hpcp.grenoble.hp.com/.

APC Uninterruptible Power Supplies for Superdome

The Superdome team has qualified the APC Silcon 3-phase 20 kW UPS for Superdome.

There are several configurations that can be utilized depending on the Superdome configuration your customer is deploying. They range from a 64 processor Superdome with dual cord and dual UPS with main tie main to a 32 processor Superdome with single cord and single UPS. In all configurations the APC Silcon SL20KFB2 has been tested and qualified by the Superdome engineers to ensure interoperability.

HP UPS Solutions

Product Number/ Description	Quantity/ Configuration	Watt	VA	Technology	Family	Package	Output
SL20KFB2 APC Silcon 3-phase UPS	• Quantity 2 32- or 64-processor dual-cord/dual-UPS with main-tie-main • Quantity 1 32- or 64-processor single-cord/single-UPS	20 kW	20 kVA	Delta conversion on-line double conversion	APC Silcon 3-phase	Standalone rack	Configurable for 200: 208 or 220V 3-phase nominal output voltage
QJB22830 Switch Gear	• Quantity 1 32- or 64-processor dual-cord/dual-UPS with main-tie-main • Quantity 0 32- or 64-processor single-cord/single-UPS	N/A	N/A	N/A	Customer Design for Superdome	N/A	N/A
WSTRUP5X8-SL10 Start-Up Service	• Quantity 2 32- or 64-processor dual-cord/dual-UPS with main-tie-main • Quantity 1 32- or 64-processor single-cord/single-UPS	N/A	N/A	N/A	N/A	N/A	N/A
WONSITENBD-SL10 Next Business Day On-site Service	• Quantity 2 32- or 64-processor dual-cord/dual-UPS with main-tie-main • Quantity 1 32- or 64-processor single-cord/single-UPS	N/A	N/A	N/A	N/A	N/A	N/A

NOTE: The APC Silcon 3-phase UPS solutions for Superdome must be ordered directly from APC. Please contact Ron Seredian at rseredia@apcc.com.

Superdome Server Watt Ratings for UPS loading

Class	Models	Watt Rating for UPS loading	UPSs Typically Used
Superdome	32-processor	19 kW	SL20KFB2; 20 kW/20 kVA
Superdome	64-processor	19 kW each cabinet; 38 kW total	SL20KFB2; 20 kW/20 kVA; Quantity 2

Configuration

Power Protection

Runtimes

The UPS will provide battery backup to allow for a graceful shutdown in the event of a power failure. Typical runtime on the APC SL20KFB2 Silcon 3 Phase UPS varies with the kW rating and the load. The APC SL20KFB2 UPS provides a typical runtime of 36.7 minutes at half load and 10.7 at full load. If additional run time is needed please contact your APC representative

Power Conditioning

The APC SL20KFB2 provides unparalleled power conditioning with its Delta-Conversion on-line double conversion technology. This is especially helpful in regions were power is unstable.

Continuous Power during Short Interruptions of Input Power

The APC SL20KFB2 will provide battery backup to allow for continuous power to the connected equipment in the event of a brief interruption in the input power to the UPS. Transaction activity will continue during brief power outage periods as long as qualified UPS units are used to provide backup power to the SPU, the Expansion Modules, and all disk and disk array products.

UPS Configuration Guidelines

In general, the sum of the "Watt rating for UPS sizing" for all of the connected equipment should not exceed the watt rating of the UPS from which they all draw power. In previous configuration guides, this variable was called the "VA rating for UPS sizing." With Unity Power Factor, the Watt rating was the same as the kVA rating, so it didn't matter which one we used. VA is calculated by multiplying the voltage times the current. Watts, which is a measurement of true power, may be less than VA if the current and voltage are not in phase. APC SL20KFB2 has Unity Power Factory correction, so the kW rating equals the kVA rating. Be sure to add in the needs for the other peripherals and connected equipment. When sizing the UPS, allow for future growth as well. If the configuration guide or data sheet of the equipment you want to protect gives a VA rating, use this as the watt rating. If the UPS does not provide enough power for the additional devices such as system console and mass storage devices, additional UPSs may be required.

Superdome

The only qualified UPS available for use with Superdome is the APC SL20KFB2 Silcon 3 Phase 20-kW UPS.
The APC SL20KFB2 can provide power protection for the SPU and peripherals. If the system console and primary mass storage devices also require power protection (which is highly recommended) they may require one or more additional UPSs depending on the total Watts. Make sure that the total watts do not exceed the UPS's voltage rating.

Integration/Installation

The APC SL20KFB2 includes both field integration start up service and next day on-site service for one year provide by APC.

Power Connections with the APC SL20KFB2

Product Number	Watts	NOM Out	Output Receptacles	Input Receptacles
SL20KFB2	20 kW	115/200 3PH, 120/208 3PH, 127/220 3PHV	Hardwire	Hardwire

Communications Connections

A DB-25 RS-232 Contact Closure connection is standard on all APC SL20KFB2 UPS. A WEB/SNMP card is also included.

Power Management

Description	Network interface cards that provide standards-based remote management of UPSs
General Features	Boot-P support, Built-in Web/SNMP management, Event logging, Flash Upgradeable, MD5 Authentication Security, Password Security, SNMP Management, Telnet Management, Web Management
Includes	CD with software, User Manual
Documentation	User Manual Installation Guide

Type of UPSs

Some customers may experience chronic "brown-out" situations or have power sources that are consistently at the lower spectrum of the standard voltage range. For example, the AC power may come in consistently at 92 VAC in a 110 VAC area. Heavy-load electrical equipment or power rationing are some of the reasons these situations arise. The APC SL20KFB2 units are designed to kick in before the AC power drops below the operating range of the HP Superdome Enterprise Server. Therefore, these UPS units may run on battery frequently if the AC power source consistently dips below the threshold voltage. This may result in frequent system shutdowns and will eventually wear out the battery. Although the on-line units can compensate for the AC power shortfall, the battery life may be shortened. The best solution is to use a good quality boost transformer to "correct" the power source before it enters the UPS unit.

Ordering Guidelines

- The APC SL20KFB2 Silcon 3-phase UPS units may be ordered as part of a new Superdome system order or as a field upgrade to an existing system.
- For new systems order please contact Ron Seredian at APC by e-mail at rseredia@apcc.com during the Superdome pre-consulting phase. APC will coordinate with HP to ensure the UPS is installed to meet the Superdome installation schedule.
- For field upgrades please contact Ron Seredian at APC by e-mail at rseredia@apcc.com when you determine a customer is in need and/or interested in power protection for Superdome. APC will coordinate with the customer to ensure the UPS is installed to meet their requirements.
- Numerous options can be ordered to compliment APC SL20KFB2 Silcon 3-phase UPS units. Your APC consultant can review these option with you are you can visit the APC website at www.apcc.com

Power Redundancy

Superdome servers, by default, provide an additional power supply for N+1 protection. As a result, Superdome servers will continue to operate in the event of a single power supply failure. The failed power supply can be replaced without taking the system down.

Multi-cabinet Configurations

When configuring Superdome systems that consist of more then one cabinet and include I/O expansion cabinets, certain guidelines must be followed, specifically the I/O interface cabling between the Superdome cabinet and the I/O expansion cabinet can only cross one additional cabinet due to cable length restrictions.

Configuration Guidelines/Rules

Superdome Configuration Guidelines/Rules

Configuration

Category	Rule Index	Rule Description
General	1	Every Superdome complex requires connectivity to a Support Management Station (SMS). The PC-based SMS also serves as the system console.
	2	Every cell in a Superdome complex must be assigned to a valid physical location.
Processor	3	All processors in a cell are the same type, same Front Side Bus (FSB) frequency, and same core frequency.
Memory	4	Configurations with 8, 16 and 32 DIMM slots are recommended (i.e. are fully qualified and offer the best bandwidth performance.)
	5	Configurations with 4 and 24 DIMM slots are supported (i.e. are fully qualified, but don't necessarily offer the best bandwidth performance).
	6	DIMMs can be deallocated in 2 DIMM increments (to support HA).
	7	Mixed DIMM sizes within a cell board are supported, but only in separate Mbat interleaving groups.
	8	System orders from the factory provide mixed DIMM sizes in recommended configurations only.
	9	For system orders from the factory, the same memory configuration must be used for all cells within a partition.
	10	DIMMs in the same rank must have SDRAMs with the same number of banks and row and column bits.
	11	Size of memory within an interleave group must be power of 2.
	12	DIMMs within the same interleave group must be same size and have same number of banks, row bits, and column bits.
	13	There are currently no restrictions on mixing DIMMs (of the same type) with different vendor SDRAMs.
I/O	14	One cell in every partition must be connected to an I/O chassis that contains a Core I/O card, a card connected to boot media, a card connected to removable media, and a network card with a connected network.
	15	A partition cannot have more I/O chassis than it has active cells.
	16	Removable media device controller should be in slot 8 of the I/O chassis.
	17	Core I/O card must be in slot 0 of the I/O chassis.(sx2000 systems do not have a core I/O)
	18	Boot device controller should be in slot 1 of the I/O chassis
	19	PCI-X high bandwidth I/O cards should be in the high bandwidth slots in the I/O chassis
	20	Every I/O card in an I/O chassis must be assigned to a valid physical location.
	21	Every I/O chassis in a Superdome complex must be assigned to a valid physical location
Performance	22	The amount of memory on a cell should be evenly divisible by 4 GB if using 512-MB DIMMs or 8 GB if using 1-GB DIMMs, i.e. 8, 16 or 32 DIMMs. The cell has four memory subsystems and each subsystem should have an echelon (2 DIMMs) populated. The loading order of the DIMMs alternates among the four subsystems. This rule provides maximum memory bandwidth on the cell, by equally populating all four memory subsystems.
	23	All cells in a partition should have the same number of processors.
	24	The number of active processors per cell should be balanced across the partition, however minor differences are OK. (Example: 4 active processors on one cell and three active processors on the second cell)
	25	If memory is going to be configured as fully interleaved, all cells in a partition should have the same amount of memory (symmetric memory loading). Asymmetrically distributed memory affects the interleaving of cache lines across the cells. Asymmetrically distributed memory can create memory regions that are non optimally interleaved. Applications whose memory pages land in memory interleaved across just one cell can see up to 16 times less bandwidth than ones whose pages are interleaved across all cells.

	26	If a partition contains 4 or fewer cells, all the cells should be linked to the same crossbar (quad) in order to eliminate bottlenecks and the sharing of crossbar bandwidth with other partitions. In each Superdome cabinet, slots 0, 1, 2 and 3 link to the same crossbar and slots 4, 5, 6 and 7 link to the same crossbar.
	27	A Core I/O card should not be selected as the main network interface to a partition. A Core I/O card is a PCI X 1X card that possibly produces lower performance than a comparable PCI X 2X card.
	28	The number of cells in a partition should be a power of two, i.e., 2, 4, 8, or 16. Optimal interleaving of memory across cells requires that the number of cells be a power of two. Building a partition that does not meet this requirement can create memory regions that are non optimally interleaved. Applications whose memory pages land in the memory that is interleaved across just one cell can experience up to 16 times less bandwidth than pages which are interleaved across all 16 cells.
	29	Before consolidating partitions in a Superdome 32 processor or 64 processor system, the following link load calculation should be performed for each link between crossbars in the proposed partition. Links loads less then 1 are best. As the link load begins to approach 2 performance bottlenecks may occur. For crossbars X and Y Link Load = Qx * Qy / Qt / L, where - Qx is the number of cells connected to crossbar X (quad) - Qy is the number of cells connected to crossbar Y (quad) - Qt is the total number of cells in the partition - L is the number of links between crossbar X and Y (2 for Superdome 32 processor systems and 1 for Superdome 64 processor systems)
	30	Maximum performance for optimal configurations (power of two cells, uniform memory across cells, power of two DIMM ranks per cell)
	31	(If rule #30 cannot be met, rule #31 is recommended) Non-power of two cells, but still uniform memory across cells, power of two DIMM ranks per cell, uniform type of DIMM.
	32	(If rule #30 or #31 cannot be met, rule #32 is recommended) Same amount of memory in each cell, but possibly different memory types in each cell (for instance, a two cell configuration with 8 512MB DIMMs (sx1000 only) in one cell, and 4 1GB DIMMs in the other). Differences in memory across different cells within the same partition should be minimal for the best performance.
	33	Same amount of memory in each cell, but non optimal and/or mixed loading within a cell (for instance, a two cell configuration with 16 512MB DIMMs (sx1000 only) and 8 1GB DIMMs in each cell).
	34	Non-uniform amount of memory across cells (this needs to boot and run, but performance is whatever you get).
	35	For the same amount of total memory, best performance is with a larger number of smaller size DIMMs.

Single System High Availability	36	Each cell should have at least two active processors.
	37	Each cell should have at least 4 GB (8 DIMMs) of memory using 512 MB DIMMs (sx1000 only) and at least 8 GB of memory using 1 GB DIMMs.
	38	I/O chassis ownership must be localized as much as possible. One way is to assign I/O chassis to partitions in sequential order starting from INSIDE the single cabinet, then out to the I/O expansion cabinet 'owned' by the single cabinet.
	39	I/O expansion cabinets can be used only when the main system cabinet holds maximum number of I/O card cages. Thus, the cabinet must first be filled with I/O card cages before using an I/O expansion cabinet.
	40	Single cabinets connected to form a dual cabinet (using flex cables) should use a single I/O expansion cabinet if possible.
	41	Spread enough connections across as many I/O chassis as it takes to become 'redundant' in I/O chassis'. In other words, if an I/O chassis fails, the remaining chassis have enough connections to keep the system up and running, or in the worst case, have the ability to reboot with the connections to peripherals and networking intact.
	42	All SCSI cards are configured in the factory as unterminated. Any auto termination is defeated. If auto termination is not defeatable by hardware, the card is not used at first release. Terminated cable would be used for connection to the first external device. In the factory and for shipment, no cables are connected to the SCSI cards. In place of the terminated cable, a terminator is placed on the cable port to provide termination until the cable is attached. This is needed to allow HP-UX to boot. The customer does not need to order the terminators for these factory integrated SCSI cards, since the customer will probably discard them. The terminators are provided in the factory by use of constraint net logic.
	43	Partitions whose I/O chassis are contained within a single cabinet have higher availability than those partitions that have their I/O chassis spread across cabinets.
	44	A partition's core I/O chassis should go in a system cabinet, not an I/O expansion cabinet
	45	A partition should be connected to at least two I/O chassis containing Core I/O cards. This implies that all partitions should be at least 2 cells in size. The lowest number cell or I/O chassis is the 'root' cell; the second lowest number cell or I/O chassis combo in the partition is the 'backup root' cell.
	46	A partition should consist of at least two cells.
	47	Not more than one partition should span a cabinet or a crossbar link. When crossbar links are shared, the partition is more at risk relative to a crossbar failure that may bring down all the cells connected to it.
Multi-System High Availability (Please also refer to Multi-System High Availability section following this table)	48	Multi-initiator support is required for Serviceguard.

Configuration

Traditional Multi-System High Availability	49	To configure a cluster with no SPOF, the membership must extend beyond a single cabinet. The cluster must be configured such that the failure of a single cabinet does not result in the failure of a majority of the nodes in the cluster. The cluster lock device must be powered independently of the cabinets containing the cluster nodes. Alternative cluster lock solution is the Quorum Service, which resides outside the Serviceguard cluster providing arbitration services.
	50	A cluster lock is required if the cluster is wholly contained within two single cabinets (i.e., two Superdome/ 16 processor or 32 processor systems or two Superdome/PA 8800 /PA 8900 32 processor or 64 processor systems) or two dual cabinets (i.e. two Superdome 64 processor systems or two Superdome/PA 8800/PA 8900 128 processor systems). This requirement is due to a possible 50% cluster failure.
	51	Serviceguard only supports cluster lock up to four nodes. Thus a two cabinet configuration is limited to four nodes (i.e., two nodes in one dual cabinet Superdome/ 64 processor system or Superdome/PA 8800/PA 8900 128 processor system and two nodes in another dual cabinet Superdome/ 64 processor system or Superdome/PA 8800/PA 8900 128 processor system). The Quorum Service can support up to 50 clusters or 100 nodes (can be arbitrator to both HP UX and Linux clusters).
	52	Two-cabinet configurations must evenly divide nodes between the cabinets (i.e. 3 and 1 is not a legal 4-node configuration).
	53	Cluster lock must be powered independently of either cabinet.
	54	Root volume mirrors must be on separate power circuits.
	55	Redundant heartbeat paths are required and can be accomplished by using either multiple heartbeat subnets or via standby interface cards.
	56	Redundant heartbeat paths should be configured in separate I/O chassis when possible.
	57	Redundant paths to storage devices used by the cluster are required and can be accomplished using either disk mirroring or via LVM's pvlinks.
	58	Redundant storage device paths should be configured in separate I/O chassis when possible.
	59	Dual power connected to independent power circuits is recommended.
Heterogeneous Multi System High Availability	60	Cluster configurations can contain a mixture of Superdome and non Superdome nodes.
	61	Care must be taken to configure an even or greater number of nodes outside of the Superdome cabinet
	62	If half the nodes of the cluster are within a Superdome cabinet, a cluster lock is required (4-node maximum cluster size)
	63	If more than half the nodes of a cluster are outside the Superdome cabinet, no cluster lock is required (16-node maximum Serviceguard cluster size).
	64	Up to a 4 node cluster is supported within a single cabinet system (Superdome/16 processor or Superdome/PA 8800/PA 8900 32 processor)
	65	Up to an 8 node cluster is supported within a single cabinet system* (Superdome/32 processor or Superdome/PA 8800/PA 8900 64 processor)
	66	Up to a 16 node cluster is supported within a dual cabinet system* (Superdome/64 processor or Superdome/PA 8800/PA 8900 128 processor)
	67	Cluster lock is required for 2-node configurations
	68	Cluster lock must be powered independently of the cabinet.
	69	Root volume mirrors must be on separate power circuits.
	70	Dual power connected to independent power circuits is highly recommended.

Configuration

* Superdome 32 processor system requires an I/O expansion cabinet for greater than 4 nodes. Superdome 64 processor system requires an I/O expansion cabinet for greater than 8 nodes.

NOTE:
"Recommended" refers to configurations that are fully qualified and offer the best bandwidth performance.
"Supported" refers to configurations that are fully qualified, but do not necessarily offer the best performance.

Instant Capacity (iCAP)-formerly known as Instant Capacity on Demand (iCOD) and Pay Per Use Programs. For a complete description of how to configure Instant Capacity, please refer to the following URL:
http://h18000.www1.hp.com/products/quickspecs/11723_div/11723_div.HTML

OpenVMS
NOTE: Support for iCAP, TiCAP, and PPU on sx1000 and sx2000 Superdome configurations is available with OpenVMS V8.3 (or higher).

Upgrades

For information on Superdome System Upgrades, please refer to the Superdome Server Upgrades QuickSpec.

Memory

Total Amount Memory per Cell	No. of 512 MB	No. of 1 GB	E0	E1	E2	E3	E4	E5	E6	E7	E8	E9	EA	EB	EC	ED	EE	EF
			0A-0B	1A-1B	2A-2B	3A-3B	4A-4B	5A-5B	6A-6B	7A-7B	8A-8B	9A-9B	AA-AB	BA-BB	CA-CB	DA-DB	EA-EB	FA-FB
4 GB	8		512 MB	512 MB	512 MB	512 MB												
8 GB		8	1 GB	1 GB	1 GB	1 GB												
8 GB	16		512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB								
16 GB		16	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB								
16 GB	32		512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB
32 GB		32	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB
12 GB	8	8	1 GB	1 GB	1 GB	1 GB	512 MB	512 MB	512 MB	512 MB								
24 GB	16	16	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB
28 GB	8	24	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	512 MB	512 MB	512 MB	512 MB
2 GB	4		512 MB	512 MB														
4 GB		4	1 GB	1 GB														
12 GB	24		512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB				
24 GB		24	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB				
20 GB	8	16	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	512 MB	512 MB	512 MB	512 MB				

Ex Echelon number, i.e. Echelon 0 consists of 2 DIMMs, 1 on A side and 1 on B side.
0A 0B refers to two DIMMs in Echelon 0, A and B side.

Recommended List of DIMM Configurations in Superdome (sx1000 systems)

Total Amount of Memory Per Cell (GB)	Number of 512 MB DIMMs	Number of 1 GB DIMMs
2	4	0
4	8	0
4	0	4
8	0	8
8	16	0
12	8	8
12	24	0
20	8	16
16	0	16
16	32	0
24	16	16
24	0	24
28	8	24
32	0	32

Memory

NOTES:

1. Configurations with 8, 16, or 32 DIMMs will result in the best performance
2. These are configurations that are shipped from manufacturing. Other configurations are supported, as long as they are not illegal.

Total Amount Memory per Cell	No. of 1 GB	No. of 2 GB	E0	E1	E2	E3	E4	E5	E6	E7	E8	E9	EA	EB	EC	ED	EE	EF
			0A-0B	1A-1B	2A-2B	3A-3B	4A-4B	5A-5B	6A-6B	7A-7B	8A-8B	9A-9B	AA-AB	BA-BB	CA-CB	DA-DB	EA-EB	FA-FB
8 GB	8		1 GB	1 GB	1 GB	1 GB												
16 GB		8	2 GB	2 GB	2 GB	2 GB												
16 GB*	16		1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB								
24 GB	8	8	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB								
24 GB	24		1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB				
32 GB*		16	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB								
32 GB	16	8	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB				
32 GB*	32		1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB
40 GB	8	16	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB				
40 GB	24	8	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB
48 GB		24	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB				
48 GB*	16	16	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB
58 GB	8	24	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB
64 GB*		32	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB

*NOTE: These configurations will result in the best performance

sx2000 Superdome Memory Configurations

Memory per Cell (GBytes)	Number of DIMMs			Echelon														
	1 GB	2 GB	4 GB	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
8	8			1 GB	1 GB	1 GB	1 GB												
16		8		2 GB	2 GB	2 GB	2 GB												
16	16			1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB								
24	8	8		2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB								
24	24			1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB				
32			8	4 GB	4 GB	4 GB	4 GB												
32	16	8		2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB				
32		16		2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB								
32	32			1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB
40	8		8	4 GB	4 GB	4 GB	4 GB	1 GB	1 GB	1 GB	1 GB								
40	8	16		2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB				
40	24	8		2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB
48		8	8	4 GB	4 GB	4 GB	4 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB				
48	16		8	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB				
48		24		2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB				
48	16	16		2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB
56	8	8	8	4 GB	4 GB	4 GB	4 GB	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB				

Memory

56	8	24		2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB
56	24		8	4 GB	4 GB	4 GB	4 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB
64			16	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB								
64		16	8	4 GB	4 GB	4 GB	4 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB				
64	16	8	8	4 GB	4 GB	4 GB	4 GB	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB
64		32		2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB
72	8		16	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	1 GB	1 GB	1 GB	1 GB				
72	8	16	8	4 GB	4 GB	4 GB	4 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB
80		8	16	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	2 GB	2 GB	2 GB	2 GB				
80		24	8	4 GB	4 GB	4 GB	4 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB
80	16		16	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB
88	8	8	16	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	2 GB	2 GB	2 GB	2 GB	1 GB	1 GB	1 GB	1 GB
96			24	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB				
96		16	16	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB	2 GB
104	8		24	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	1 GB	1 GB	1 GB	1 GB
112		8	24	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	2 GB	2 GB	2 GB	2 GB
128			32	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB	4 GB

*NOTE: These configurations will result in the best performance

Superdome sx2000 Specifications	SPU Description	Superdome 16 processor*	Superdome 32 processor*	Superdome 64 processor*
	SPU Product Number	A9833A	A9834A	A9834A+A9835A
	Number of Dual core Itanium 2 1.6 GHz processors	1 - 16	1 - 32	2 - 64 Windows Server 2003, DatacenterEnterprise Edition can support a maximum of 32 Dual core 8 Itanium2 processors with 64 cores per partition if hyper-threading is turned off; or a maximum of 16 Dual core Itanuim2 processors with 32 cores and 64 threads per partition if hyper-threading is turned on. Windows Server 2003, Enterprise Edition can support a maximum of 8 Dual core Itanium2 processors with 16 cores, and 32 threads per partition (see ordering guidelines under the operating system section).
	Number of Itanium 2 1.6 GHz processors	2-16	2-32	4-64 Windows Server 2003, Enterprise Edition can support a maximum of 8 Itanium2 processors per partition (see ordering guidelines under the operating system section).
	Benchmarks			
	TPC C disclosure (HP UX)	N/A	N/A	TBD
	TPC C disclosure (Windows)	N/A	N/A	N/A
	Supported Processors			
	Dual core Itanium 2 processor	1.6 GHz, 24 MB & 18MB cache	1.6 GHz, 24 MB & 18MB cache	1.6 GHz, 24 MB & 18MB cache
	Itanium 2 Processor	1.6 GHz, 9 MB cache	1.6 GHz, 9 MB cache	1.6 GHz, 9 MB cache
	Memory			
	1 GB DIMMs (DDRII)	8-128 GB	8-128 GB	16-512 GB
	2 GB DIMMs (DDRII)	16-256 GB	16-256 GB	32-1024 GB

4-GB DIMMs (DDRII)	32-512 GB	32-1024 GB	64-2048 GB (max. is 1024 GB for Windows)
DIMM Density (MB)	1024/2048/4096	1024/2048/4096	1024/2048/4096
Cell Boards			
2 processor 4 processor cells	1 - 4	1 - 8	4 - 16
Software Support			
HP-UX revision	HP UX 11i version 2		
Windows revision	Windows Server 2003, Datacenter Edition for Itanium 2 . Hyper-threading supported with Dual core Itanium2 processors (default from factory hyper-threading off).		
Linux	RHEL 4 U3 and SLES 10		
OpenVMS	OpenVMS V8.3 minimum; dual core (Montecito) processors only on sx2000 configurations; hyperthreading support on sx2000		
Expandability / Connectivity			
12-slot PCI-X I/O chassis SPU cabinet must be filled first before placing I/O chassis in I/O expansion cabinet	1 - 4 No I/O expansion cabinet needed	1-8 I/O expansion cabinet required if number of I/O chassis is more greater than 4.	1-16 I/O expansion cabinet required if number of I/O chassis is greater than 8. A second I/O expansion cabinet is required if the number of I/O chassis is greater than 14. NOTE: Linux to support up to 8 I/O chassis by Q2 2007
Number of partitions without I/O expansion cabinet	1 - 4	1 - 4	1 - 8
Number of partitions with I/O expansion cabinet	Not Applicable	1 - 8	1- 16
RS-232C serial ports	Yes	Yes	Yes
10/100Base-T Ethernet	Yes Windows does not support 10/100Base-T Ethernet. Windows requires Gigabit Ethernet card	Yes Windows does not support 10/100Base-T Ethernet. Windows requires Gigabit Ethernet card	Yes Windows does not support 10/100Base-T Ethernet. Windows requires Gigabit Ethernet card

Superdome sx1000 Specifications				
	SPU Description	**Superdome 16 processor***	**Superdome 32 processor***	**Superdome 64 processor***
	SPU Product Number	A6113A	A5201A	A5201A+A5202A

Number of Itanium 2 1.5 GHz or 1.6 GHz processors	2 - 16	2 - 32	6 - 64 Windows Server 2003 Enterprise Edition supports a maximum of 8 processors per partition (see ordering guidelines under the operating system section).
Number of mx2 processors NOTE: Linux and OpenVMS do not support mx2	2 - 32	2 - 64	6 - 128 Windows Server 2003 Datacenter Edition supports up to 64 processors in a single instance of the operating system. A fully configured mx2 128 processor Superdome requires at least two instances of Windows Server 2003 Datacenter Edition. Windows Server 2003 Enterprise Edition is not supported with mx2 modules
Benchmarks			
TPC-C disclosure (HP-UX)	N/A	N/A	TBD
TPC-C disclosure (Windows)	N/A	N/A	1,231,433 tpmC Windows Server 2003 Datacenter Edition with SQL Server 2005 64 bit version
Supported Processors			
Itanium 2 Processor	1.5 GHz, 6 MB cache	1.5 GHz, 6 MB cache	1.5 GHz, 6 MB cache
Itanium 2 Processor	1.6 GHz, 9 MB cache	1.6 GHz, 9 MB cache	1.6 GHz, 9 MB cache
Mx2 Processor Module (2 processor cores)	1.1 GHz, 6 MB cache (each processor)	1.1 GHz, 6 MB cache (each processor)	1.1 GHz, 6 MB cache (each processor)
Memory			
512 MB DIMMs (sx1000 only)	2 - 128 GB	2 - 256 GB	6 - 256 GB
1 GB DIMMs	4 - 128 GB	4 - 256 GB	12 - 512 GB
2 GB DIMMs	8 - 256 GB	8 - 512 GB	24 - 1024 GB
DIMM Density (MB)	512 / 1024	512 / 1024	512 / 1024
Cell Boards			

1 processor (mx2 only)/2 processor 3 processor (mx2 only)/4 processor cells	1 - 4	1 - 8	3 - 16

Software Support

HP-UX revision	HP- UX 11i version 2		
Windows revision	Windows Server 2003, Datacenter Edition for Itanium 2; Windows Server 2003, Enterprise Edition for Itanium 2		
Linux revision	Red Hat RHEL AS 3,and 4 /SUSE SLES 9 and 10 Linux revision on Superdome with Intel Itanium 2 processors only (not on Superdome with mx2 processor modules)		
OpenVMS revision	OpenVMS version 8.2-1 (or higher) Superdome with Intel Itanium 2 Mad9M processors only (not on Superdome with mx2 or Mad6 modules)		

Expandability / Connectivity

12 slot PCI X I/O chassis Note: SPU cabinet must be filled first before placing I/O chassis in I/O expansion cabinet	1 - 4 No I/O expansion cabinet needed	1 - 8 I/O expansion cabinet required if number of I/O chassis is more greater than 4. Linux supports 2 I/O Chassis	3 - 16 I/O expansion cabinet required if number of I/O chassis is greater than 8. A second I/O expansion cabinet is required if the number of I/O chassis is greater than 14.
Number of Partitions without I/O expansion cabinet	1 - 4	1 - 4	1 - 8
Number of Partitions with I/O expansion cabinet	Not Applicable	1 - 8	1- 16
RS 232C Serial Ports	Yes	Yes	Yes
10/100Base T Ethernet	Yes Windows does not support 10/100 Base T Ethernet. Windows requires Gigabit Ethernet card	Yes Windows does not support 10/100 Base T Ethernet. Windows requires Gigabit Ethernet card	Yes Windows does not support 10/100 Base T Ethernet. Windows requires Gigabit Ethernet card

Maximum I/O Cards See supported I/O table for specific products

SPU Product Number	A6113A(sx1000) A9833A (sx2000)	A5201A (sx1000) A9834A (sx2000)	A5201A+A5202A (sx1000) A9834A+A9835A (sx2000)
Mass Storage	16 - 48	16 - 96	16 - 192
LAN	2 - 32	2 - 48	2 - 96
WAN	8	16	32
Multi-Function (Mass Storage / LAN)	8 - 48	16 -96	32 -192

Additional Interface Cards	8	8 -14	8 -14
Typical Electrical Characteristics by Configuration			
AC input power	200 240 VAC, 50/60 Hz	200 240 VAC, 50/60 Hz	200 240 VAC, 50/60 Hz
Typical Input Power (watts) dissipation Itanium 2 (watts)	5,900 4 cells, 32 GB, 4 I/O chassis with 6 PCI each	9,800 8 cells, 32 GB, 4 I/O chassis with 6 PCI each	19,600 16 cells, 32 GB, 4 I/O chassis with 6 PCI each
Typical power dissipation mx2 (watts)	5,730	9,490	18,980
Typical 3 Phase 4Wire Current requirements at 200V 240V	17A / Phase	29A / Phase	Two Cabinets 29A / Phase
Typical 3 Phase 5Wire Current requirements at 200V 240V	10A / Phase	17A / Phase	Two Cabinets 17A / Phase
Maximum Electrical Characteristics			
Maximum Input Power (VA)	8,200	12,196	24,392
Cabinet Marked Electrical			
AC input power-Option 7: 3 phase 5 wire input	200 240 VAC phase to neutral, 5 wire, 50/60 Hz		
AC input power-Option 6: 3 phase 4 wire input	200 240 VAC phase to phase, 4 wire, 50/60 Hz		
Maximum Current requirements at 220V 240V:			
Option 7: 3 phase 5 wire input	24 A/Phase	24 A/Phase	24 A/Phase
Option 6: 3 phase 4 wire input	44 A/Phase	44 A/Phase	44 A/Phase
Required Power Receptacle-	Option 6: IEC Compliant 4 Wire 60A 250V Ph-Ph Option 5: IEC Compliant 5 Wire 32A 480VPh-Ph		
Heat Characteristics			
Maximum Heat dissipation (BTUs/hour)	28,969	41,614	83,288
Typical Heat dissipation (BTUs/hour)	20,131	33,439	66,877
Site Preparation			
Site planning and installation included	Yes	Yes	Yes
Depth (mm / inches)	1,220 / 48.1	1,220 / 48.1	1,220 / 48.1
Width (mm / inches)	762 / 30	762 / 30	1,524 / 60
Height (mm / inches)	1,960 / 77.2	1,960 / 77.2	1,960 / 77.2
Weight (Kg / lbs)	500 / 1,123	598 / 1,343	1,196 / 2,685

Environmental Characteristics

Acoustics	65 dB
Operating temperature	+20° to +30° C
Non operating temperature	-40° to +70° C
Maximum rate of temperature change	20° C/hour
Operating relative humidity	15% to 80% @ 30°C
Operating altitude	0 3.1 km (0 10,000 ft)
Non operating altitude	0 4.6 km (0 15,000 ft)

Regulatory Compliance

Safety	IEC 950:1991+A1, +A2, +A3, +A4; EN60950:1992+A1, +A2, +A3, +A4, +A11; UL 1950, 3rd edition; cUL CSA C22.2 No. 950 95

Key Dates

First CPL date	6/03 (sx1000) 3/06 (sx2000)
First ship date	3Q03 (sx1000) 2Q06 (sx2000)

Given that Itanium 2 1.5 GHz are single core processors and mx2 is a dual core processor, the columns listed in this table refer to 16 processor, 32 processor and 64 processor. This terminology refers to 16 core, 32 core and 64 core for Superdome Itanium 2 1.5 GHz systems and 32 processor cores, 64 processor cores and 128 processor cores for Superdome mx2 systems.

sx1000 and sx2000 Superdome Specifications

SPU Description	Superdome 16 processor*	Superdome 32 processor*	Superdome 64 processor*
SPU Product Number	A6113A (sx1000) A9833A (sx2000)	A5201A (sx1000) A9834A (sx2000)	A5201A+A5202A (sx1000) A9834A+A9835A (sx2000)
Maximum I/O Cards			
Mass storage	16-48	16-96	16-192
LAN	2-32	2-48	2-96
WAN	8	16	32
Multi function (Mass Storage/LAN)	8-48	16-96	32-192
Additional interface cards	8	8-14	8-14
Typical Electrical Characteristics by Configuration			
AC input power	200-240 VAC, 50/60 Hz	200-240 VAC, 50/60 Hz	200-240 VAC, 50/60 Hz
Typical input power dissipation Itanium 2 (watts)	5,900 4 cells, 32 GB, 4 I/O chassis with 6 PCI each	9,800 8 cells, 32 GB, 4 I/O chassis with 6 PCI each	19,600 16 cells, 32 GB, 4 I/O chassis with 6 PCI each
Typical power dissipation mx2 (watts)	5,730	9,490	18,980

Typical 3 phase 4 wire current requirements at 200V 240V	17A/Phase	29A/Phase	Two cabinets 29A/Phase
Typical 3 phase 5 wire current requirements at 200V 240V	10A/Phase	17A/Phase	Two cabinets 17A/Phase
Maximum Electrical Characteristics			
Maximum input power (VA) cabinet marked electrical	8,200	12,196	24,392
AC input power-Option 7: 3 phase 5 wire input	200 240 VAC phase to neutral, 5 wire, 50/60 Hz		
AC input power-Option 6: 3 phase 4 wire input	200 240 VAC phase to phase, 4 wire, 50/60 Hz		
Maximum current requirements at 220V 240V:			
Option 7: 3 phase 5 wire input	24 A/Phase	24 A/Phase	24 A/Phase
Option 6: 3 phase 4 wire input	44 A/Phase	44 A/Phase	44 A/Phase
Required power receptacle	Option 6: IEC Compliant 4 wire 60A 250V Ph Ph Option 5: IEC Compliant 5 wire 32A 480V Ph Ph		
Heat Characteristics			
Maximum heat dissipation (BTUs/hour)	28,969	41,614	83,288
Typical heat dissipation (BTUs/hour)	20,131	33,439	66,877
Site Preparation			
Site planning and installation included	Yes	Yes	Yes
Depth (mm/inches)	1,220/48.1	1,220/48.1	1,220/48.1
Width (mm/inches)	762/30	762/30	1,524/60
Height (mm/inches)	1,960/77.2	1,960/77.2	1,960/77.2
Weight (Kg/lbs)	500/1,123	598/1,343	1,196/2,685
Environmental Characteristics			
Acoustics	65 dB		
Operating temperature	+20°C to +30°C		
Non operating temperature	-40°C to +70°C		
Maximum rate of temperature change	20°C/hour		
Operating relative humidity	15% to 80% @ 30°C		
Operating altitude	0 3.1 km (0 10,000 ft)		
Non operating altitude	0 4.6 km (0 15,000 ft)		

Regulatory Compliance

Safety	IEC 950:1991+A1, +A2, +A3, +A4; EN60950:1992+A1, +A2, +A3, +A4, +A11; UL 1950, 3rd edition; cUL CSA C22.2 No. 950 95

Key Dates

First CPL date	6/03 (sx1000); 3/06 (sx2000)
First ship date	3Q03 (sx1000); 2Q06 (sx2000)

*NOTE: Given that Itanium 2 1.5 GHz is a single core processor and mx2 is a dual core processor, the columns listed in this table refer to 16 processor, 32 processor and 64 processor. This terminology refers to 16 processor cores, 32 processor cores and 64 processor cores for Superdome Itanium 2 1.5 GHz systems and 32 processor cores, 64 processor cores and 128 processor cores for Superdome mx2 systems.

Superdome I/O Expansion (IOX) Cabinet Specifications

SPU Description	Superdome 16 processor	Superdome 32 processor	Superdome 64 processor
SPU Product Number	A6113A (sx1000) A9833A (sx2000)	A5201A (sx1000) A9834A (sx2000)	A5201A+A5202A (sx1000) A9834A+A9835A (sx2000)
Maximum Number of I/O Chassis Enclosures (ICEs)*	Not Applicable	2	4
Peripherals Supported	Not Applicable	All peripherals qualified for use with Superdome and/or for use in a 10K G2 or Rack System E are supported in the I/O expansion cabinet as long as there is available space. Peripherals not connected to or associated with the Superdome system to which the I/O expansion cabinet is attached may be installed in the I/O expansion cabinet.	
Servers Supported	No servers except those required for Superdome system or High Availability Observatory or ISEE may be installed in an I/O expansion cabinet.		
Dimensions			
Height (mm / inches)		1.6 meters or 2.0 meters	
Depth (mm / inches)	Not Applicable	45.5 in (115.57 cm) (same depth as 32W)	
Width (mm / inches)		24.0 in (61 cm)	
Environmental			
	Not Applicable	Same as Superdome	
Relevant Product Numbers			

12 slot PCI X chassis for HP rack expansion cabinet	Not Applicable	A6864AZ (sx1000) A9836AZ (sx2000)	A6864AZ (sx1000) A9836AZ (sx2000)
I/O Expansion Cabinet Power and Utilities Subsystem		A5861A	A5861A
I/O Expansion Power and Utilities Subsystem Graphite color		A5861D	A5861D
I/O Chassis Enclosure for 12 slot PCI X Chassis		A5862A (sx1000) A9852A (sx2000)	A5862A (sx1000) A9852A (sx2000)

NOTE: Each ICE holds two I/O card cages or 24 PCI X I/O slots.

APC SL20KFB2 Specifications

Description		APC Silcon 20000VA/20000W Input: 115/200 3PH, 120/208 3PH, 127/220 3PHV; Output: 115/200 3PH, 120/208 3PH, 127/220 3PHV; Interface port: DB 25 RS 232; Contact closure
General features		0 95% non condensing, 200% overload capability, audible alarms, built in static bypass switch, Delta Conversion online technology, environmental protection, event logging, extendable run time, full rated output available in kW, input power factor correction, intelligent battery management, LCD alphanumeric display, overload indicator, paralleling capability, sine wave output, SmartSlot, software, web management
Includes		Parallel card, triple chassis for three SmartSlots, user manual, Web/SNMP Management Card
Spare parts kits		See APC website http://www.apcc.com
Documentation		User Manual and Installation Guide
Input	Nominal input voltage	115/200 3PH, 120/208 3PH, 127/220 3PH V
	Input frequency	50 Hz programmable +/- 0.5, 1, 2, 4, 6, 8%; 60 Hz programmable +/- 0.5, 1, 2, 4, 6, 8%
	Input connection type	Hardwire 5 wire (3PH+N+G)
	Input voltage range for main operations	170-230 (200 V), 177-239 (208 V), 187-242 (220 V) V
Batteries	Typical backup time at half load	36.7 minutes
	Typical backup time at full load	10.7 minutes
	Battery type	Maintenance free sealed lead acid battery with suspended electrolyte: leak proof
	Typical recharge time[1]	2 hours
Physical	Maximum height dimensions	55.12 in (140.00 cm)
	Maximum width dimensions	39.37 in (100.00 cm)
	Maximum depth dimensions	31.50 in (80.01 cm)
	Net weight	1,290.00 lbs (586.36 kg)
	Shipping weight	1,340.00 lbs (609.09 kg)

	Shipping height	66.93 in (170.00 cm)
	Shipping width	43.31 in (110.00 cm)
	Shipping depth	35.43 in(90.00 cm)
	Color	Dark green (NCS 7020 B50G), Light gray (NCS 2703 G84Y)
	Units per pallet	1.0
	Communications and Management	
Communications and Management	Interface port	DB-25 RS-232, Contact Closure
	SmartSlot interface quantity	2
	Pre installed SmartSlot cards	AP9606
	Control panel	Multi-function LCD status and control console
	Audible alarm	Beep for each 52 alarm conditions
	Emergency power off (EPO)	Yes
	Optional management device	See APC website http://www.apcc.com
Environmental	Operating environment	32° to 104°F (0° to 40 °C)
	Operating relative humidity	0% to 95%
	Operating elevation	0 to 3333 ft (0 to 999.9 m)
	Storage temperature	-58° to 104°F (-50° to 40 °C)
	Storage relative humidity	0% to 95%
	Storage elevation	0 to 50,000 ft (0 to 15,000 m)
	Audible noise at 1 meter from surface of unit	55 dBA
	Protection class	4,094 BTU/hour NEMA 1, NEMA 12 optional
	Online thermal dissipation	4,094 BTU/hour
	Protection Class	NEMA 1, NEMA 12
Conformance	Approvals	EN 55022 Class A, ISO 9001, ISO 14001, UL 1778, UL Listed, cUL Listed
	Standard warranty	One-year repair or replace, optional on-site warranties available, optional extended warranties available
	Optional new service	See APC website http://www.apcc.com

[1]The time to recharge to 90% of full battery capacity following a discharge to shutdown using a load rated for 1/2 the full load rating of the UPS

Sx2000 Superdome supported I/O

I/O Card	Product Number	Connector Type(s)	Operating Systems[2]	Maximum Cards by System (per partition)			Special Notes
				16	32	64	

Mass Storage Host Bus Adapters							
PCI 2x Fibre Channel	A5158A	LC	H	48	96	192	
PCI 2 Gb/s Fibre Channel	A6795A	LC	H	48	96	192	
PCI-X 2 channel 2 Gb /s Fibre Channel	A6826A	LC	H, O	48	96	192	OpenVMS: Max 8/partition
PCI 1 port U160 SCSI	A6828A		H	48	96	192	
PCI 2 port U160 SCSI	A6829A		H	48	96	192	
PCI 1 port 4Gb FC card (PCIX-266)	AB378A	LC	H, O	48	96	192	OpenVMS: Max 8/partition
PCI 1 port 4Gb FC card (PCIX 266)	AB378B	LC	H	48	96	192	
PCI 2 port 4Gb FC card (PCIX 266)	AB379A	LC	H , W, L, O	48	96	192	Windows: Maximum 16/partition; No boot support from factory; boot support capable with complete system integration in the field) OpenVMS and Linux: Max 8/partition
PCI 2 port 4Gb FC card (PCIX 266)	AB379B	LC	H	48	96	192	
PCI-X 4 Gb Fibre Channel Adapter FC1114BR	AB429A		W, L	24	24	24	Windows: No boot support from factory; boot support capable with complete system integration in the field) Linux: 8 per partition
PCI 2 channel Ultra320 SCSI	A7173A	VHDCI	H, W, L, O	48	96	192	Windows: 16 per partition, no boot support OpenVMS: Max 4/partition; bootable Linux: 8 per partition
PCI-X 2 channel Smart Array 6402 U320	A9890A	VHDCI	W, L	8	8	8	Boot support
PCI-X 4-channel Smart Array 6404 Ultra320	A9891A	VHDCI	W	8	8	8	Windows: Boot support

PCI-X 4Gb Fiber Channel Adapter FC2143BR	AD167A		W, L	24	24	24	Windows: No boot support from factory; boot support capable with complete system integration in the field) Linux: 8 per partition
PCI-X 2P 4Gb Fiber Channel Adapter FC2243BR	AD168A	LC	W, L	16	16	16	Windows: No boot support from factory; boot support capable with complete system integration in the field) Linux: 8 per partition
PCI-X 4 Gb Fibre Channel Adapter FC1114BR	AB429A		W, L	24	24	24	Windows: No boot support from factory; boot support capable with complete system integration in the field) Linux: 8 per partition
Smart Array P600 serial attached SCSI (SAS) Controller - for external storage connect only (Maximum cable length for P600/MSA50 connect is 4 meters).	337972-B21	2 internal (SFF8484) x 4 wide and 1 external (SFF8470) x 4 wide	W	8	8	8	Windows: Boot support
512MB cache memory upgrade for SA P600 and SA6402 Controller (field install only)	372538-B21		W				
PCI X 2 channel 2 Gb /sFibre Channel Windows	AB466A	LC	W	16	16	16	Windows: No boot support from factory; boot support capable with complete system integration in the field)

PCI X 1 channel 2 Gb /sFibre Channel Windows	AB467A	LC	W	32	32	32	Windows: No boot support from factory; boot support capable with complete system integration in the field)
Local Area Network (LAN) Adapters							
PCI 1 port 1000Base SX	A4926A	RJ 45	H	16	32	64	
PCI 1 port 1000BaseT	A4929A	RJ-45	H	16	32	64	
PCI 1 port 1000BaseT	A5230A	RJ-45	H	24	48	96	
PCI 4 port 100Base TX	A5506B	RJ-45	H	8	16	32	
PCI 1 port 1000Base T (gigabit copper)	A6825A	RJ-45	H, O	16	32	64	OpenVMS: Maximum 8/partition; bootable
PCI 1 port 1000Base SX (gigabit fiber)	A6847A	Duplex SC	H, O	16	32	64	OpenVMS: Maximum 8/partition; boot support
PCI 1 port 1000Base SX (gigabit fiber)	AD332A	Duplex SC	H, O	16	32	64	
PCI-X 2-port 1000Base-SX	A7011A	Duplex SC	H, O	16	32	64	OpenVMS: Maximum 8/partition; boot support
PCI-X 2-port 1000Base-T	A7012A	RJ-45	H, O	16	32	64	OpenVMS: Maximum 8/partition; boot support
PCI 1 port 1000Base-T	A7061A	RJ 45	W	32	32	32	Windows: No boot support
PCI 1 port 1000Base-SX	A7073A	Duplex SC	W	32	32	32	Windows: No boot support
PCI 2 port Windows/Linux 1000Base SX	A9899A	LC	W, L	16	16	16	Windows: Boot support Linux: 8 per partition
PCI 2 port Windows/Linux 1000Base TX	A9900A	RJ-45	W, L	16	16	16	Windows: Boot support Linux: 8 per partition
PCI -X 2 port 4x Fabric (HPC) Adapter	AB286A	4x Infiniband Copper	H	8	8	8	
HP PCI-X 4-port 1000Base-T Gigabit Adpt	AB545A	RJ-45	H, O	16	32	64	OpenVMS: Maximum 3/partition; boot support

HP PCI-X 2-port 4X Fabric (HA & DB) Adpt	AB345A	4x Infiniband Copper	H	16	32	64	No boot support
HP PCI-X 10 GigE	AB287A	Duplex LC	H	8[1]	8[1]	16[1]	
PCI-X 10GbE LAN Adapter	AD144A	Duplex LC	W	8	8	8	Windows: no boot support from factory, boot support capable with complete system integration in the field Linux: 2 per partition
NC370T PCI-X MFN 10/100/1000T (copper)	374191-B22		H	16	32	64	
NC370F PCI-X MF 1000SX (fibre)	374193-B22		H	16	32	64	
Multi-Function Cards (Mass Storage / LAN)							
PCI 2 port 100Base T/ 2 port Ultra2 SCSI	A5838A	VHDCI/RJ-45	H[1]	8	16	32	
PCI-X 2Gb Fibre Channel / 1000BaseSX	A9782A	LC	H, V	48	96	192	
PCI-X 2Gb Fibre Channel / 1000BaseTX	A9784A	1 LC, 1 RJ-45	H	48	96	192	
HP PCI-X Multifunction 2-port 2Gb FC / 2-port 1 Gb Ethernet Adapter	AB465A	2 LC	H	48	96	192	
HP PCI-X Multifunction 2-port 1000BT and dual-port U320 SCSI adapter	AB290A	SCSI - LVD/SE LAN - RJ-45	H	48	96	192	
Graphics/USB card	A6869B		H, W, L, O	1	1	1	Windows and Linux: Supports graphics and USB, no boot support Linux and OpenVMS: USB only
Wide Area Network (WAN) Adapters							
2 port Programmable Serial Interface (PSI) X.25 / Frame Relay / SDLC	J3525A	RS-530, RS-232, V.35, RS-449 or X.21	H	8	16	32	
Additional Interface Cards							
PCI-X 2-port HBA 4x Fabric (HA & DB)	AB345C		H	8	16	32	
PCI-X 2-port HBA 4x Fabric (HPC)	AB286C		H	8	16	32	
PCI HyperFabric 2 Fibre[1]	A6386A	LC Duplex	H	8	8	8	
PCI 8-port Terminal Multiplexer	A6748A	RS-232	H	8	14	14	

PCI 64-port Terminal Multiplexer	A6749A	RS-232 or RS-422	H	8	14	14	

[1] Maximum cards are two per I/O chassis. Therefore: 16 processor=8 maximum, 32 processor=16 maximum, and 64 processor=32 maximum.
[2] Operating Systems Legend H=HP UX 11i V2, W= Windows Server 2003 Datacenter or Enterprise Edition, L= Linux (Red Hat RHELAS 3 and 4 and SuSE SLES9), O= OpenVMS

sx1000 Superdome Supported I/O

I/O Card	Product Number	Connector Type(s)	Operating Systems[2]	Maximum Cards by System (per partition)			Special Notes
				16	32	64	
Mass Storage Host Bus Adapters							
PCI 2x Fibre Channel	A5158A	LC	H	48	96	192	
PCI 2 Gb/s Fibre Channel	A6795A	LC	H	48	96	192	
PCI X 2 channel 2 Gb /sFibre Channel	A6826A	LC	H, L, O	48	96	192	14 for SLES 9, 10 8 for RHEL 4 OpenVMS-8 per partition; bootable
PCI 1 port 4Gb FC card (PCIX 266)	AB378A		O	48	96	192	OpenVMS: Max 8/partition ; V8.3 minimum
PCI 2 port 4 Gb FC card (PCIX 266)	AB379A		O	48	96	192	OpenVMS: Maximum 8/partition; V8.3 minimum Linux : 8 per partition
PCI 2 channel Ultra320 SCSI	A7173A	VHDCI	H, W, L,O	48	96	192	Windows: 16 per partition, no boot support OpenVMS: 4 per partition, bootable Linux : 8 per partition
PCI-X 2 channel Smart Array 6402 U320	A9890A	VHDCI	W , L				Windows: Boot support 8/Linux partition
PCI-X 4-channel Smart Array 6404 Ultra320	A9891A	VHDCI	W	8	8	8	Windows: Boot support
Smart Array P600 serial attached SCSI (SAS) Controller (for external storage connect only) (Maximum cable length for P600/MSA50 connect is 4 meters).	337972-B21	2 internal (SFF8484) x 4 wide and 1 external (SFF8470) x 4 wide	W	8	8	8	Windows: Boot support
512MB cache memory upgrade for SA P600 and SA6402 Controller (field install only)	372538-B21		W				

PCI-X 4Gb Fiber Channel Adapter FC2143BR	AD167A		W	24	24	24	Windows: No boot support from factory; boot support capable with complete system integration in the field)
PCI X 2 channel 2 Gb /sFibre Channel Windows	AB466A	LC	W	16	16	16	(Windows: No boot support from factory; boot support capable with complete system integration in the field)
PCI X 1 channel 2 Gb /sFibre Channel Windows	AB467A	LC	W	32	32	32	Windows: No boot support from factory; boot support capable with complete system integration in the field)
Local Area Network (LAN) Adapters							
PCI-X 4-Gb Fiber Channel Adapter FC2143BR (Windows maximum=24 per partition/system)	AD167A	LC	W, L	24 (W)	24 (W)	24 (W)	Windows: No boot support from factory; boot support capable with complete system integration in the field) Linux : 8 per partition
PCI-X 2P 4-Gb Fiber Channel Adapter FC2243BR (Windows maximum=16 per partition/system)	AD168A	LC	W, L	16 (W)	16 (W)	16 (W)	Windows: No boot support from factory; boot support capable with complete system integration in the field) Linux : 8 per partition
PCI 1 port Universal FDDI LAN	A3739B	FDDI SC					
PCI 1 port 1000Base SX	A4926A	RJ 45	H	16	32	64	
PCI 1 port 1000BaseT	A4929A	RJ-45	H	16	32	64	
PCI 1 port 1000BaseT	A5230A	RJ-45	H	24	48	96	
PCI 4 port 100Base TX	A5506B	RJ-45	H, L	8	16	32	

PCI 1 port 802.5 Token Ring 4/16/100	A5783A	RJ-45 and DB 9	H	8	16	32	
PCI 1 port 1000Base T (gigabit copper)	A6825A	RJ-45	H, O	16	32	64	OpenVMS - 8/part
PCI 1 port 1000Base SX (gigabit fiber)	A6847A	Duplex SC	H, O	16	32	64	OpenVMS - 8/part
PCI-X 2-port 1000Base-SX	A7011A	Duplex SC	H	16	32	64	
PCI X 2-port 1000Base-SX	A7011A	Duplex LC	O				OpenVMS - 8/part
PCI-X 2-port 1000Base-T	A7012A	RJ-45	H, O	16	32	64	OpenVMS - 8/part
PCI 1 port 1000Base-T	A7061A	RJ 45	W, L	32	32	32	8 for Linux Windows: No boot support
PCI 1 port 1000Base-SX	A7073A	Duplex SC	W, L	32	32	32	8 for Linux Windows: No boot support
PCI 2 port Windows/Linux 1000Base SX	A9899A	LC	W,L	16	16	16	Windows: Boot support Linux : 8 per partition
PCI 2 port Windows/Linux 1000Base TX	A9900A	RJ-45	W,L	16	16	16	Windows: Boot support Linux : 8 per partition
PCI -X 2 port 4x Fabric (HPC) Adapter	AB286A	4x Infiniband Copper	H	8	8	8	
HP PCI-X 4-port 1000 Base-T Gigabit Adapter	AB545A	RJ-45	H, O	16	32	64	Stand alone shipments OpenVMS: 3 per partition, no boot support
HP Linux 4-port 1000Base-T Gigabit Adptr	AD145A	RJ45	L	8			Boot support Can be used only with RHEL 4 U1 or later 4 cards per partition
HP PCI-X 2-port 4X Fabric (HA & DB) Adapter	AB345A	4x Infiniband Copper	H	8	8	8	Stand-alone Shipments
HP PCI-X 4-port 1000Base-T Gigabit Adpt	AB545A	RJ-45	H, O	16	32	64	Boot support OpenVMS - 3/part; no boot support
HP PCI-X 2-port 4X Fabric (HA & DB) Adpt	AB345A	4x Infiniband Copper	H	16	32	64	No boot support
HP PCI-X 10 GigE	AB287A	Duplex LC	H	8	8	16	No boot support

HP Linux 133MHz 10GbE SR Fiber Adapter	AD144A	Duplex LC	L, W	4 8 (W)	8 (W)	8 (W)	Windows: No boot support from factory. Boot support capable with complete system integration in the field; 8 per partition Boot support RHEL 4 U1 or greater, SLES 9 SP3 or SLES 10 Linux: 2 cards per partition
Multi-Function Cards (Mass Storage / LAN)							
PCI 2 port 100Base T/ 2 port Ultra2 SCSI	A5838A	VHDCI/RJ-45	H[1]	8	16	32	
PCI-X 2Gb Fibre Channel / 1000BaseSX	A9782A	LC	H, O	48	96	192	OpenVMS - 4/part; FC bootable
PCI-X 2Gb Fibre Channel / 1000BaseTX	A9784A	1 LC, 1 RJ-45	H, O	48	96	192	OpenVMS - 4/part; FC bootable
HP PCI-X Multifunction 2-port 2Gb FC / 2-port 1 Gb Ethernet Adapter	AB465A	2 LC	H, O	48	96	192	• 3.3 volts only • Supported on HP UX 11i and 11i v2 • Boot support on FC and 1-Gb Ethernet • OpenVMS-2/partition; FC bootable
HP PCI-X Multifunction 2-port 1000BT and dual-port U320 SCSI adapter	AB290A	SCSI - LVD/SE LAN - RJ-45	H, O	48	96	192	Boot support on SCSI and GigE OpenVMS - 2/part; SCSI bootable
HP PCI X Multi function 2 port FC and 2 port GigE-T	AB465A	2 LC, 2 RJ-45	H, O	48	96	192	Boot support on FC and GigE OpenVMS-2/partition; FC bootable

Wide Area Network (WAN) Adapters							
PCI 1 port ATM 155 Mbps (MMF)	A5513A	Duplex SC	H	8	16	32	
2 port Programmable Serial Interface (PSI) X.25 / Frame Relay / SDLC	J3525A	RS-530, RS-232, V.35, RS-449 or X.21	H	8	16	32	
Additional Interface Cards							
PCI HyperFabric 2 Fibre[1]	A5486A	LC Duplex	H	8	8	8	
PCI 8-port Terminal Multiplexer	A6748A	RS-232	H	8	14	14	
PCI 64-port Terminal Multiplexer	A6749A	RS-232 or RS-422	H	8	14	14	
HP StorageWorks Linux Q2300 64-bit HBA	A7538A		L	8	8	8	
1-port 4Gb FC Qlogic	AB429A		L	8	8	8	RHEL 4 U4 and SLES 10
Standalone Switches							
Description	**Product Number**	**Connector**	**Form Factor**	**Operating Systems**			**Notes**
HP 24 port 4X Fabric Copper Switch	376227-B21		1U	H			Replaces AB399A
PCI X 24 port 4x Fabric Copper Switch	AB399A	4x Infiniband Copper	1U	H			Discontinued, Replaced by 376227 B21

[1]No Serviceguard support
[2]Operating Systems Legend H=HP-UX 11i V2, W= Windows Server 2003 Datacenter or Enterprise Edition, L= Linux (Red Hat RHELAS 3 and SUSE SLES9), O= OpenVMS
[3]Supported, but may no longer be orderable
[4]Standalone orders only

NOTE: MSA30 SB/DB are supported as boot disks for Superdome running HP UX 11i with the following cards: A7173A, A6828A, A6829A and A5838A.

Technical Specifications

The boot configuration for Windows Server 2003, Datacenter Edition Superdome partitions can be:

- The Smart Array 6402/6404 disk array controllers (A9890A/A9891A) connected to StorageWorks 4400 (a.k.a. MSA30) series enclosures.
- The SA P600 SAS controller (337972 B2) connected to the MSA50 (NOTE: Maximum cable length available is 4 meters.)
- Fibre Channel PCI X (AB466A or AB467A) HBAs(no boot from factory).

The Windows Server 2003 operating system comes with a software mirroring solution. However, the majority of Windows customers use hardware based RAID solutions (such as the industry leading Smart Array disk array controllers from HP) instead and do not use this mirroring tool. Also note that the Smart Array controllers do not support failover capability (customers cannot have 2 Smart Array cards connected to the same boot partition on a StorageWorks 4300/4400 enclosure). RAID levels 0, 1, 5, 1+0 and ADG are supported as well as disk sparing.

Note that booting from external storage arrays is supported (HP XP and EVA storage). In these cases, it is recommended by HP that the FC HBAs are configured in a redundant pair using HP Secure Path software for high availability.

Please refer to the table below as guidance for configuring your **Windows Server 2003 partition** on Superdome (note that "Watson" rules are in place that reflect these recommendations). Please note that if the VGA/USB card (A6869A/A6869B) is used, it would only be needed once per instance to the Windows OS instance.

PCI-X Technical Slotting Information for Windows Server 2003

	Left								Right			
Slot	11	10	9	8	7	6	5	4	3	2	1	0
Clock Speed (MHz)	66	66	66	66	66 or 133	66 or 133	66 or 133	66 or 133	66	66	66	66
Special Notes for Windows Server 2003 Datacenter Edition				SCSI Card (A7060A Or A7173A) removable media slot					Default boot device slot for Smart Array controller (A9890A) recommended	Windows LAN Gig E card (A7061A)	VGA/USB (A6869A/ A6869B). If ordered, must be placed in Slot One	Core I/O slot (A6865A)
Previously on A4856A	2X	2X	2X	2X	4X	4X	4X	4X	2X	2X	2X	2X
Now on PCI-X	4X	4X	4X	4X	8X	8X	8X	8X	4X	4X	4X	4X

NOTE: FC HBA (AB466A or AB467A) are to consume 8X slots first and then populate 4X slots (recommended for performance optimization)

To ensure Windows Server 2003 high availability for storage connectivity, it is recommended to use HP SecurePath (with HP storage) and EMC PowerPath (with EMC storage) for load balancing/redundancy between fibre channel HBAs (AB232A, AB466A or AB467A).

External Storage

HP has the broadest, most robust server and storage line up in the industry, providing exactly the right fit for every need. Refer to the Storage-Server matrix found in the introductory section of Chapter 4 of this configuration guide to see a matrix that highlights which storage device, server and operating system is interoperable.

For HP Storage connectivity, the webpage http://spock.corp.hp.com/index.aspx has all of the detail for HP hardware. Please consult this matrix for HP supported on line and near line storage.

For EMC connectivity with Windows Server 2003 on HP Integrity servers, the EMC support matrix has detailed information concerning supported HP hardware: http://www.emc.com/interoperability/index.jsp. Please consult this matrix to determine if your customer's desired configuration is supported by EMC.

Superdome Supported Online Storage

Storage Device	HP-UX 11i v2	Windows Server 2003 Datacenter Edition	Red Hat RHEL AS 3 &4 and SUSE SLES 9	OpenVMS
XP 48/512	Yes	Yes	Yes	Yes
XP128/1024	Yes	Yes	Yes	Yes
VA7100	Yes	No	Yes	No
VA7400	Yes	No	Yes	No
VA 7410/7110	Yes	Yes	Yes	No
MSA1000	Yes	Yes	Yes	Yes
MSA1500	No	Yes	No	No
EVA 5000	Yes	Yes (EVA v3 or greater)	Yes	Yes
EVA 3000	Yes	Yes (EVA v3 or greater)	Yes	Yes
StorageWorks 4300 series	No	Yes	Yes	Yes
StorageWorks 4400 series (MSA30)	Yes	Yes	Yes	Yes
MSA50	No	Yes	No	No
FC10	Yes	No	No	No
SC10	Yes	No	No	No
DS2100	Yes	Yes (not recommended for use with U320 adapter)	Yes	Yes
DS2120	No	Yes (not recommended for use with U320 adapter)	No	No
DS2110	Yes	Yes (not recommended for use with U320 adapter)	Yes	Yes
DS2300	Yes	No	Yes	No
DS2405	Yes	No	Yes	No
EMC Symmetrix 3000	Yes	No	No	Yes
EMC Symmetrix 5000	Yes	No	No	Yes
EMC Symmetrix 5500	Yes	No	No	Yes
EMC Symmetrix 8000	Yes	Yes	No	Yes
EMC DMX Series	Yes	Yes	No	Yes
EMC CLARiiON CX200	No	Yes	No	No
EMC CLARiiON CX 400/CX600	No	Yes	No	No
EMC CLARiiON CX300/CX500/CX700		Yes		No
EMC CLARiiON FC4700	No	Yes	No	No
SAN 2/8	Yes	No	Yes	Yes

SAN 2/8 EL	Yes	No	Yes	Yes
SAN 2/16	Yes	No	Yes	Yes
SAN 2/16 EL	Yes	No	Yes	Yes
StorageWorks Core 2/64	Yes	Yes	Yes	Yes
StorageWorks Edge 2/24	Yes	No	Yes	Yes
StorageWorks Edge 2/32	Yes	No	Yes	Yes
StorageWorks SAN Director 2/64	Yes	Yes	Yes	Yes
StorageWorks SAN Director 2/140	Yes	No	Yes	Yes

Superdome Supported Nearline Storage

Storage Device	HP-UX 11i v2	Windows Server 2003 Datacenter Edition	Red Hat RHEL AS 3 & 4 and SUSE SLES 9	OpenVMS
ESL9595 with SDLT 220 and 320	Yes	Yes	Yes	Yes
ESL9595 with Ultrium 230 and 460 drives	Yes	Yes	Yes	No (230) / Yes (460)
ESL9322 with SDLT 220 and 320	Yes	Yes	Yes	Yes
ESL9322 with Ultrium 230 and 460 drives	Yes	Yes	Yes	No (230) / Yes (460)
MSL5000 series with Ultrium 230 drives	Yes	Yes	Yes	No
MSL5000 series with SDLT 220 drives	Yes	Yes	Yes	Yes
MSL5000 series with SDLT 320 drives	Yes	Yes	Yes	Yes
MSL6000 series with Ultrium 460 drives	Yes	Yes	Yes	Yes
SSL1016 with DLT1	Yes	Yes	Yes	No
SSL1016 with SDLT 320	Yes	Yes	Yes	No
SSL1016 with Ultrium 460	Yes	Yes	Yes	No
Tape Autoloader 1/8	Yes	Yes	Yes	Yes
NSR 1200 FC/SCSI router for MSL series libraries	Yes	No	Yes	Yes
NSR e1200, e1200-160 FC/SCSI router for MSL libraries	Yes	No	Yes	Yes
NSR e2400, e2400-160 FC/SCSI router for ESL libraries	Yes	No	Yes	Yes
NSR 2402 FC/SCSI router for ESL series libraries	Yes	No	Yes	Yes
Optical Jukebox 2200mx	Yes	No	No	No
Optical Jukebox 1200mx	Yes	No	No	No
Optical Jukebox 700mx	Yes	No	No	No
Optical Jukebox 600mx	Yes	No	No	No
Optical Jukebox 300mx	Yes	No	No	No
Optical Jukebox 220mx	Yes	No	No	No
Optical Jukebox 9100mx	Yes	No	No	No
Ultrium 460 Standalone/Rack	Yes	Yes	Yes	Yes
Ultrium 230 Standalone/Rack	Yes	Yes	Yes	No
Ultrium 215 Standalone/Rack	Yes	Yes	Yes	No

DVD - Rack	Yes	Yes	Yes	No
TA5300 - Tape Array (plus all supported devices in TA5300)*	Yes	Yes	Yes	Yes
DDS-4 Standalone/Rack	Yes	Yes	Yes	Yes
DDS-4×6 Standalone	Yes	No	Yes	Yes
DDS-5 Standalone/Rack	Yes	Yes	Yes	Yes
DLT-80 Standalone/Rack	Yes	No	Yes	No
DLTVS80 Standalone/Rack	Yes	No	Yes	No

*Supported devices in TA5300 are C7498A, C7497B, Q1524B, C7492B, DW019A, C7470B, Q1512B, Q1540A

NOTES:

- All shipments of SCSI devices for Superdome except HVD10 and SC10 are supported with standard cables and auto termination enabled. Only the Surestore Disk System HVD10 (A5616AZ) and the Surestore Disk System SC10 (A5272AZ) will use disabled auto termination and In-Line Terminator cables.
- Each A5838A PCI 2-port 100Base-T 2-port Ultra2 SCSI card that supports a Surestore Disk System SC10 (A5272AZ) will need quantity two (2) of product number C2370A (terminator); otherwise it must have a terminated cable in place prior to HP UX boot.

Peripherals Required Per Partition (nPar)

	HP-UX 11i v2	Windows Server 2003	Red Hat RHEL AS 3 & 4 and SUSE SLES 9	OpenVMS
I/O Cards	• Core I/O (Slot 0) provides console and LAN • Default Boot Device (Slot 1) • Removable Media Card (Slot 8)	• Core I/O (Slot 0) provides console only, Windows does not support the 10/100 LAN • A7061A, A7073A, A9899A or A9900A provide LAN support (Slot 2) • Optional A6869A/ A6869B Obsidian Card or Obsidian 2 Card (Slot 1)- USB/VGA • Removable Media Card A7059A/A7060A/ A7073A (Slot 8)	• Core I/O (Slot 0) provides console and LAN • Default Boot Device (Slot 1) • Removable Media Card A7059A/A7060A (Slot 8)	• Core I/O (Slot 0) provides console and LAN • Boot Device (any slot) • Removable Media Card (any slot)
Peripherals	• DVD Hard Drive (Boot Disk) • DDS-4/DAT-40 Tape Backup • C7508AZ or C7508A (Qualec Device)	• DVD Hard Drive (Boot Disk) • DDS-4/DAT-40 Tape Backup • C7508AZ or C7508A Tape Array 5300	• DVD Hard Drive (Boot Disk) • DDS-4/DAT-40 Tape Backup • C7508AZ or C7508A (Tape Array 5300)	• DVD Hard Drive